



Delft University of Technology

**Document Version**

Final published version

**Citation (APA)**

Ethembaraoglu, A. M. (2026). *Patchwork security: Municipal Cybersecurity Measures in Practice*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:bccd21fa-b328-44d4-8405-cbb4c72fd88c>

**Important note**

To cite this publication, please use the final published version (if applicable).  
Please check the document version above.

**Copyright**

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.  
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

**Sharing and reuse**

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

**Takedown policy**

Please contact us and provide details if you believe this document breaches copyrights.  
We will remove access to the work immediately and investigate your claim.

*This work is downloaded from Delft University of Technology.*

# PATCHWORK SECURITY

Municipal Cybersecurity Measures in Practice



Aksel M. Ethembaoglu



**PATCHWORK SECURITY**  
**MUNICIPAL CYBERSECURITY MEASURES IN PRACTICE**



**PATCHWORK SECURITY**  
**MUNICIPAL CYBERSECURITY MEASURES IN PRACTICE**

**Dissertation**

for the purpose of obtaining the degree of doctor  
at Delft University of Technology,  
by the authority of the Rector Magnificus, Prof. dr. ir. H.Bijl,  
chair of the Board for Doctorates  
to be defended publicly on  
Thursday, 2 July 2026, 15:00

by

**Aksel Mahir ETHEMBABA OGLU**

This dissertation has been approved by the promotors and the copromotor.

Prof. dr. M.J.G. van Eeten  
Dr. R.S. van Wegberg  
Dr. Y. Zhauniarovich

Composition of the doctoral committee:

Rector Magnificus  
Prof. dr. M.J.G. van Eeten  
Dr. R.S. van Wegberg  
Dr. Y. Zhauniarovich

Chairperson  
Delft University of Technology, promotor  
Delft University of Technology, promotor  
Delft University of Technology, copromotor

*Independent members:*

Prof.dr.ing. A.J. Klievink  
Prof.dr. T.A.P. Metze  
Dr. L. Allodi  
Dr. I. Westerman  
Prof.dr. M.E. Warnier

University of Leiden  
Delft University of Technology  
Eindhoven University of Technology  
Ministry of Interior  
Delft University of Technology, *reserve member*

This research was supported by the Ministry of the Interior and Kingdom Relations of the Netherlands and Delft University of Technology under Grant M75B07.



*Keywords:* vulnerability scanning, patching, vulnerability notifications, municipalities, local government, CERT, CSIRT, threat intelligence, APT, attribution  
*Printed by:* Gildeprint  
*Cover image:* Wendy Bour

Copyright © 2026 by A.M. Ethembabaoglu

ISBN: 978-94-6518-336-7

An electronic version of this dissertation is available at  
<http://repository.tudelft.nl/>

*For Michelle. And for my sons Ezra and Elias, who were born during this PhD.  
Love is always the answer.*



# CONTENTS

<b>Summary</b>	<b>xi</b>
<b>Samenvatting</b>	<b>xv</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Background . . . . .	2
1.1.1 Security, Frameworks and Compliance. . . . .	2
1.1.2 Analytical Lens for Municipal Security: Causal Model . . . . .	4
1.1.3 Preventive and Reactive Security Measures . . . . .	5
1.1.4 Surface Exposure: Managing Municipal Assets. . . . .	8
1.1.5 Threat: Know Your Attacker . . . . .	8
1.2 Research gaps. . . . .	9
1.3 Research aims and questions . . . . .	11
1.4 Dissertation outline. . . . .	11
<b>2 The Unpatchables: Why Municipalities Persist in Running Vulnerable Hosts</b>	<b>15</b>
2.1 Introduction . . . . .	16
2.2 Related Work . . . . .	18
2.3 Ethics . . . . .	19
2.4 Measurement Approach . . . . .	20
2.4.1 Scanning Municipal Networks . . . . .	20
2.4.2 User Study . . . . .	21
2.5 Validating Observed Vulnerable Systems . . . . .	24
2.6 Attribution and Unclear Responsibilities . . . . .	25
2.6.1 Role of CERT and Municipal IPs . . . . .	26
2.6.2 Misaligned Threat Landscape . . . . .	26
2.6.3 Quantifying the Misalignment . . . . .	28
2.6.4 Perspectives on Responsibilities . . . . .	28
2.7 Managing Vulnerable Systems . . . . .	30
2.7.1 Identifying Vulnerable Systems . . . . .	31
2.7.2 Vulnerable vs. No Vulnerable Systems . . . . .	31
2.7.3 No Patch . . . . .	32
2.7.4 Patching Systems: Prioritizing . . . . .	34
2.8 Discussion . . . . .	36
2.9 Conclusion . . . . .	38

<b>3</b>	<b>“Tell Them They Are a Responsible Entity, Not a Customer”: Understanding Practitioner Challenges in Sector CSIRTs</b>	<b>39</b>
3.1	Introduction . . . . .	40
3.2	Related Work . . . . .	42
3.3	Methodology . . . . .	44
3.3.1	Interview Studies . . . . .	45
3.3.2	Data and Methodology for Vulnerability Notification Analysis . . . . .	48
3.3.3	Validation Workshop . . . . .	49
3.4	Results: Challenges and Stakeholder Experiences with Sector CSIRT Services . . . . .	49
3.4.1	Incident Response: Clashing Expectations . . . . .	50
3.4.2	Advisories: Contested Value . . . . .	51
3.4.3	Expert Insights: One-Size-Fits-None . . . . .	52
3.4.4	Vulnerability Notifications: The Unseen Service . . . . .	53
3.4.5	Intelligence Sharing: Visibility Equals Authority . . . . .	56
3.4.6	Outreach and Community: Bringing Constituents Together . . . . .	56
3.5	Results: Evaluating the Vulnerability Notification Service . . . . .	57
3.5.1	Measuring Notification Flow . . . . .	58
3.5.2	Reflections Evaluating the Notification Program . . . . .	59
3.6	Results: Strategic Challenges for Sector CSIRT Practitioners . . . . .	61
3.6.1	Governance Structure and Stakeholders . . . . .	61
3.6.2	Infrastructure and Capability Management . . . . .	63
3.7	Results: Findings Validation Workshop . . . . .	63
3.8	Discussion . . . . .	65
<b>4</b>	<b>APT to Disagree: A Comparative Analysis of Attribution in Commercial TI</b>	<b>71</b>
4.1	Introduction . . . . .	72
4.2	Related Work . . . . .	73
4.3	Background . . . . .	75
4.4	Methodology . . . . .	76
4.4.1	Normalization of Heterogeneous Feeds . . . . .	76
4.4.2	Actor Name Mapping . . . . .	78
4.4.3	Agreement Metrics and Analysis . . . . .	79
4.5	TAG Validation and Augmentation . . . . .	80
4.5.1	Actor Ambiguity in TAG . . . . .	80
4.5.2	Validating TAG . . . . .	81
4.5.3	TAG Augmentation . . . . .	81
4.6	Results: Actor Tracking . . . . .	82
4.6.1	Tracked Actors Over Time . . . . .	82
4.6.2	Overlap and Union in Tracked Actors . . . . .	83
4.6.3	IOC Overlap for Jointly Tracked Actors . . . . .	85
4.7	Results: Attribution Agreement . . . . .	85
4.7.1	Co-observed IOCs . . . . .	86
4.7.2	Disagreement on MD5 Indicators . . . . .	88
4.7.3	Pairwise Agreement Among Vendors . . . . .	89

4.8	Discussion . . . . .	89
4.8.1	Implications . . . . .	90
4.8.2	Recommendations. . . . .	91
4.8.3	Limitations. . . . .	92
4.8.4	Future Research . . . . .	93
4.9	Conclusions. . . . .	93
<b>5</b>	<b>Conclusion</b>	<b>99</b>
5.1	Empirical findings . . . . .	99
5.1.1	Chapter 2 - Patch Behavior and Management of Vulnerable Hosts . . . . .	99
5.1.2	Chapter 3 - Functioning of Institutional Support Structures for Municipal Incident Prevention and Mitigation . . . . .	100
5.1.3	Chapter 4 - Quality Evaluation of Attribution in Commercial TI . . . . .	100
5.2	Reflections on Findings . . . . .	101
5.3	Governance implications . . . . .	105
5.3.1	Hierarchy . . . . .	106
5.3.2	Market . . . . .	108
5.3.3	Network . . . . .	109
5.4	Future work. . . . .	110
	<b>Bibliography</b>	<b>113</b>
<b>A</b>	<b>Appendix for Chapter 2</b>	<b>143</b>
A.1	Respondent Details . . . . .	143
A.2	Interview Protocol . . . . .	143
A.3	Codes . . . . .	145
A.4	Venn Diagrams Total IP Sets. . . . .	145
<b>B</b>	<b>Appendix for Chapter 3</b>	<b>147</b>
B.1	Interview Protocol . . . . .	147
<b>C</b>	<b>Appendix for Chapter 4</b>	<b>151</b>
C.1	Charts and Tables. . . . .	151
	<b>Acknowledgements</b>	<b>157</b>
	<b>Authorship Contributions</b>	<b>161</b>
	<b>List of Publications</b>	<b>163</b>
	<b>About the Author</b>	<b>165</b>



# SUMMARY

Municipalities play a central role in delivering essential public services, including civil registration, social services, taxation, communication, and local democratic processes. In doing so, they increasingly rely on digital systems. Cyber incidents affecting these systems can disrupt service delivery, expose sensitive personal data, and impose significant recovery costs. Because municipalities are often the most visible and accessible layer of government for citizens, such incidents may also affect public trust. In addition, municipalities operate and oversee systems that support local critical infrastructure, such as water management, traffic control, and energy distribution, placing them within the scope of both financially motivated cybercriminals and state-sponsored advanced persistent threats (APTs).

In response to this threat landscape, municipalities are expected to implement a range of cybersecurity measures. These include complying with security frameworks and standards, managing vulnerabilities through patching and configuration, participating in information sharing and coordination structures, and preparing for incident response and recovery. At the same time, municipalities typically operate under constraints that distinguish them from many other organizations, including limited internal cybersecurity capacity, extensive reliance on outsourcing and shared service providers, and complex internal structures in which responsibility for systems and data is distributed across departments and external parties.

As a result, municipal cybersecurity is rarely a matter of isolated technical controls. Instead, it is shaped by interactions between municipalities and a broader ecosystem of actors, including vendors, managed service providers, sectoral and national CSIRTs, and commercial security firms. Information about threats and vulnerabilities often reaches municipalities through intermediaries, and the ability to act on that information depends on institutional arrangements, contractual relationships, and organizational processes. Understanding municipal cybersecurity, therefore, requires examining not only which security measures are in place but also how those measures function in practice within this institutional context.

This dissertation examines the security measures municipalities use to address cyber threats and how they function in practice under these conditions. It investigates vulnerability remediation, institutional support for incident prevention and response, and the use of commercial threat intelligence, and asks how these security measures can be improved in practice, addressing the central research question: *How can municipalities improve security measures to address cyber threats?* To answer this question, the dissertation presents three empirical studies that combine technical measurements with practitioner perspectives, adopting a socio-technical approach that connects technical observations to organizational and institutional contexts.

Chapter 2, *The Unpatchables*, examines why municipalities continue to expose Internet-facing systems with known vulnerabilities for which patches exist. Using Internet-wide

measurements to identify vulnerable systems across Dutch municipalities and interviews with municipal security professionals, the chapter assesses whether such detections are false positives and investigates why remediation does not occur. The findings show that the detections were accurate, but that vulnerabilities persist largely due to misalignments between externally attributed systems and internal responsibility structures. Vulnerable systems frequently fall into gaps created by shadow IT, outdated registrations, unclear ownership, or limited ability to act. The chapter shows that vulnerability remediation depends primarily on asset management and organizational coordination rather than solely on technical feasibility.

Chapter 3 studies the functioning of institutional support structures for municipal cybersecurity, with a focus on sectoral CSIRTs. Through interviews with practitioners connected to the Dutch municipal sector CSIRT, complemented by interviews and a validation workshop with other sector CSIRT professionals, the chapter examines how services such as vulnerability notifications, incident response support, and information sharing are provided and experienced. The findings reveal persistent tensions between expectations and practice, shaped by constrained resources, ambiguous mandates, and dependencies on other actors. The chapter also identifies a breakdown in the vulnerability notification chain: notifications that should have reached the sector CSIRT were not forwarded, and the absence of feedback mechanisms prevented detection of this failure.

Chapter 4, *APT to Disagree*, examines the use of commercial threat intelligence as a security measure to support decision-making, with particular attention to threat actor attribution. Using a longitudinal analysis of attribution claims across multiple vendors and millions of indicators of compromise, the chapter evaluates both the scope of actor tracking and the degree of agreement between vendors. The results show that vendors track only a limited subset of known actors, that overlap in observed indicators is rare, and that actor-level attribution frequently disagrees even when vendors observe the same activity. While agreement is much higher at the country level, attribution remains unstable. The chapter shows that aggregating threat intelligence does not necessarily reduce uncertainty and that attribution claims may convey a misleading sense of certainty when their limitations are not explicitly communicated.

Finally, Chapter 5 synthesizes the empirical findings from Chapters 2 to 4 to answer the main research question and reflects on their implications for improving municipal cybersecurity. The chapter identifies recurring findings across the studies, including the central role of asset management and ownership in enabling effective security measures, the dependence of preventive and responsive services on institutional coordination and feedback mechanisms, and the risks associated with security information that obscures uncertainty. It reflects on municipal cybersecurity as an organizational challenge embedded in broader institutional arrangements rather than as a purely technical problem. The chapter then discusses governance implications by interpreting the findings through three modes of coordination: hierarchy, market, and network, and shows how different security measures rely on different coordination mechanisms to function effectively. Lastly, Chapter 5 outlines avenues for future work, including interdisciplinary research into asset management and vulnerability notifications, empirical studies of mitigation and recovery after compromise, cross-country comparisons of municipal and institutional arrangements, further research into CSIRT services and feedback loops, and

deeper investigation into the causes and evaluation of instability in threat actor attribution.

This dissertation shows that improving municipal security measures requires attention to how those measures operate in practice within complex institutional ecosystems. By empirically examining vulnerability remediation, sectoral support structures, and commercial threat intelligence, it provides evidence of where municipal security measures break down and identifies conditions under which they are more likely to be effective.



# SAMENVATTING

Gemeenten zijn essentieel voor de uitvoering van diverse publieke taken, zoals burgerzaken, sociale dienstverlening, belastingheffing, communicatie en lokale democratische besluitvorming. Voor deze taken zijn gemeenten in toenemende mate afhankelijk van digitale systemen. Cyberincidenten die deze systemen treffen kunnen de dienstverlening verstoren, resulteren in datalekken van gevoelige persoonsgegevens en hoge herstelkosten met zich meebrengen. Daarbij zijn gemeenten voor burgers een zichtbare vertegenwoordiging van de overheid, en kunnen cyberincidenten het vertrouwen in de overheid in het algemeen ondermijnen. Verder beheren en ondersteunen gemeenten systemen die onderdeel zijn van de lokale vitale infrastructuur, zoals waterbeheer, verkeerssystemen en energievoorziening. Daarmee zijn zij een interessant doelwit voor cyberaanvallen voor zowel financieel gemotiveerde actoren als voor statelijke actoren (advanced persistent threats, APT's).

Om deze risico's te mitigeren wordt van gemeenten verwacht dat zij uiteenlopende beveiligingsmaatregelen nemen. Het gaat hierbij om het naleven van beveiligingsstandaarden en -kaders, het verhelpen van kwetsbaarheden zoals via patching, deelname aan gremia voor samenwerkingen in informatieuitwisselingen, en het voorbereiden op incidenten, en de afhandeling en het herstel daarvan. Tegelijkertijd verschillen gemeenten op een aantal wezenlijke punten van veel andere organisaties. Zo moet er aan verschillende overheidsrandvoorwaarden worden voldaan met betrekking tot IT, beschikken zij structureel over beperkte interne cybersecuritycapaciteit, zijn zij sterk afhankelijk van uitbestede en gedeelde IT-dienstverlening en kennen zij complexe organisatorische structuren waarin verantwoordelijkheden voor systemen en gegevens zijn verdeeld over meerdere afdelingen en externe partijen.

Gemeentelijke cybersecurity kan daarom niet worden gezien als een optelsom van afzonderlijke technische maatregelen. In de praktijk wordt zij vormgegeven door een wisselwerking met een breder ecosysteem van cyber actoren, waaronder leveranciers, managed service providers, sectorale en nationale CSIRTs en commerciële beveiligingsbedrijven. Informatie over dreigingen en kwetsbaarheden bereikt gemeenten doorgaans via tussenliggende partijen, zoals CSIRTs. De mogelijkheid om op die informatie te handelen wordt in grote mate bepaald door institutionele verhoudingen, contractuele afspraken en organisatorische processen. Inzicht in gemeentelijke cybersecurity is daarom meer dan enkel de compliance maatregelen die zijn ingevoerd. Het gaat vooral om de wijze waarop deze maatregelen in de praktijk functioneren binnen deze institutionele context.

Dit proefschrift onderzoekt welke beveiligingsmaatregelen gemeenten inzetten om cyberdreigingen het hoofd te hielden en hoe deze maatregelen in de praktijk functioneren. Daarbij wordt ingezoomd op het verhelpen van kwetsbaarheden, op institutionele ondersteuning voor preventie en incidentrespons en op het gebruik van commerciële threat intelligence. De centrale onderzoeksvraag luidt: Hoe kunnen gemeenten hun

beveiligingsmaatregelen verbeteren om cyberdreigingen te mitigeren? Om deze vraag te beantwoorden presenteert dit proefschrift drie empirische studies waarin technische metingen worden gecombineerd met inzichten van professionals. Daarbij wordt een sociotechnische benadering gehanteerd, waarin technische bevindingen worden geplaatst in een organisatorische en institutionele context.

Hoofdstuk 2, *The Unpatchables*, onderzoekt waarom gemeenten systemen, die via het internet benaderbaar zijn, blijven draaien met bekende kwetsbaarheden waarvoor al patches beschikbaar zijn. Aan de hand van grootschalige internetmetingen naar kwetsbare systemen bij Nederlandse gemeenten en interviews met gemeentelijke securityprofessionals wordt gekeken of deze waarnemingen berusten op 'false positives' en waarom patches niet worden uitgerold. De resultaten laten zien dat de detecties van kwetsbare systemen kloppen, maar dat kwetsbare systemen vooral blijven bestaan door een gebrekkige aansluiting tussen extern aan gemeenten toegeschreven systemen en interne verantwoordelijkheidsstructuren. Kwetsbare systemen resideren regelmatig in blinde vlekken van een organisatie, bijvoorbeeld als gevolg van 'shadow-IT', verouderde registraties, onduidelijk systeem eigenaarschap of een beperkte mogelijkheid om daadwerkelijk in te grijpen. Het hoofdstuk laat zien dat het verhelpen van kwetsbaarheden in de praktijk vooral afhankelijk is van goed assetmanagement en organisatorische afstemming, en niet primair van technische haalbaarheid.

Hoofdstuk 3 richt zich op het functioneren van institutionele ondersteuningsstructuren voor gemeentelijke cybersecurity, in het bijzonder op sectorale CSIRTs. Op basis van interviews met professionals die betrokken zijn bij het sectorale CSIRT voor Nederlandse gemeenten, aangevuld met interviews en een validatieworkshop met medewerkers van andere sectorale CSIRTs, wordt onderzocht hoe diensten zoals kwetsbaarheidsmeldingen, ondersteuning bij incidentafhandeling en informatie-uitwisseling worden vormgegeven en ervaren. De bevindingen laten structurele spanningen zien tussen verwachtingen en de feitelijke praktijk. Deze spanningen hangen samen met beperkte capaciteit, onduidelijke mandaten en afhankelijkheden van andere organisaties. Daarnaast wordt een structurele verstoring zichtbaar in de keten van kwetsbaarheidsmeldingen: meldingen die het sectorale CSIRT hadden moeten bereiken, werden niet doorgezet, terwijl het ontbreken van terugkoppeling maakte dat deze tekortkoming niet werd gesignaleerd.

Hoofdstuk 4, *APT to Disagree*, onderzoekt het gebruik van commerciële threat intelligence (TI) als hulpmiddel voor besluitvorming, vooral voor het toeschrijven van risico's van de activiteiten aan dreigingsactoren. Op basis van een longitudinale analyse van attributieclaims van meerdere TI leveranciers en miljoenen "indicators of compromise" (IOCs) wordt zowel de omvang van actortracking als de mate van overeenstemming tussen leveranciers geanalyseerd. De resultaten tonen dat individuele TI leveranciers slechts een beperkt deel van de bekende dreigingsactoren volgen, dat overlap in waargenomen indicatoren weinig voorkomt, en dat attributie van IOCs op actorniveau vaak uiteenloopt, zelfs wanneer TI leveranciers dezelfde IOCs waarnemen. De overeenstemming tussen TI leveranciers is echter groter bij attributie op landenniveau. Het hoofdstuk laat zien dat het combineren van meerdere bronnen van threat intelligence niet vanzelfsprekend leidt tot minder onzekerheid en dat attributieclaims voor IOCs een schijnzekerheid kunnen wekken wanneer deze beperkingen niet expliciet worden gemaakt.

In Hoofdstuk 5 worden de empirische bevindingen uit Hoofdstukken 2 tot en met 4 bijeengebracht om de centrale onderzoeksvraag te beantwoorden en om te reflecteren op de betekenis van deze resultaten voor de verbetering van gemeentelijke cybersecurity. Het hoofdstuk identificeert terugkerende patronen, waaronder het belang van assetmanagement en duidelijk eigenaarschap voor het effectief kunnen uitvoeren van beveiligingsmaatregelen, de afhankelijkheid van preventieve en responsieve diensten van institutionele afstemming en feedbackmechanismen, en de risico's van beveiligingsinformatie van TI leveranciers die onzekerheid verhult. Gemeentelijke cybersecurity wordt daarmee gekarakteriseerd als een organisatorische opgave die is ingebed in een bredere institutionele context, en niet als een louter technisch vraagstuk. Vervolgens worden de governance-implicaties besproken aan de hand van drie coördinatiemechanismen — hiërarchie, markt en netwerk — en wordt inzichtelijk gemaakt hoe verschillende beveiligingsmaatregelen afhankelijk zijn van verschillende vormen van coördinatie om effectief te kunnen functioneren. Tot slot worden richtingen voor toekomstig onderzoek geschetst, waaronder interdisciplinair onderzoek naar assetmanagement en kwetsbaarheidsmeldingen, empirische studies naar mitigatie en herstel na compromittering, internationale vergelijkingen van gemeentelijke en institutionele arrangementen, verder onderzoek naar CSIRT-diensten en feedbackmechanismen, en verdiepend onderzoek naar de oorzaken en beoordeling van instabiliteit in dreigingsactorattributie.

Dit proefschrift laat zien dat het verbeteren van gemeentelijke beveiligingsmaatregelen vraagt om aandacht voor de manier waarop deze maatregelen in de praktijk functioneren binnen complexe institutionele ecosystemen. Door empirisch onderzoek te doen naar kwetsbaarheidsremediatie, sectorale ondersteuningsstructuren en commerciële threat intelligence wordt inzicht verkregen in waar gemeentelijke beveiligingsmaatregelen tekortschieten en onder welke voorwaarden zij in de praktijk effectiever zijn.



# 1

## INTRODUCTION

Municipalities worldwide are increasingly targeted by cyberattacks [100, 177]. In the United States, malware attacks on local governments rose by 148% and ransomware incidents by 51% between January and August 2023 compared to the same period in 2022 [30]. The Center for Internet Security (CIS) reported a 313% rise in endpoint security incidents in that timeframe, while Emsisoft tracked ransomware incidents affecting U.S. government entities, increasing from 95 in 2023 to 117 in 2024 [45]. Across Europe, ENISA’s Threat Landscape 2025 found that public administration was the sector reporting the most incidents (38% of all incidents), with municipalities accounting for roughly one-third of that activity [71]. These developments demonstrate that local governments have become a preferred target in the global cyber threat landscape [4, 85, 178], a trend expected to intensify as their reliance on digital systems continues to grow [119, 233].

Cyberattacks on municipalities carry high financial and operational costs for the victim municipal organizations. In the United States, major incidents have required extensive recovery spending—Atlanta (USD 17 million) [55, 207], Baltimore (USD 10 million) [87], New Orleans (USD 5.2 million) [273], and Dallas (USD 8.5 million) [42]. European municipalities have faced similar burdens: Redcar & Cleveland Borough Council (GBP 10 million) [275], Hackney Council (GBP 12 million) [46], and Hof van Twente in the Netherlands (EUR 4 million) [258]. Beyond the direct financial damage, such incidents have disrupted critical municipal functions, including email, payment, and court systems, as well as services such as housing, welfare, and civil status registration. Some attacks have even produced physical consequences: Johannesburg’s street lighting systems were frozen in the “on” position [208], while Kalix in Sweden experienced failures in heating and ventilation systems[199].

Contrary to other organizations, municipalities occupy a unique position in public life: they are the most tangible and close-to-home representation of government for citizens. They deliver essential services such as justice administration, housing, benefits, permits, garbage disposal, utilities, and identity management. When cyber incidents disable these systems or expose personal information, citizens directly experience the consequences. Such disruptions can erode trust in local government [86]—a trust that

research consistently shows to be higher than in national institutions [110]. Studies have long linked the quality of public service delivery to citizens' trust in government performance [111, 137, 253]. Persistent cyberattacks therefore not only undermine service continuity but also weaken the broader legitimacy of democratic institutions.

At the same time, municipalities often manage local critical infrastructure, including water, sewage, traffic, and energy systems, which attracts attention from state-sponsored advanced persistent threats (APTs). These actors pursue espionage or disruption rather than financial gain. The U.S. Cybersecurity and Infrastructure Security Agency reported in 2020 that a Russian-linked APT ('Berserk Bear') had targeted state and local government networks [40]. The FBI later warned of APTs exploiting municipal webserver vulnerabilities [103], and in 2023, the Municipal Water Authority of Aliquippa was attacked by the Iranian-linked group Cyber Av3ngers [39, 225, 250]. In Europe, the Sandworm group deployed Caddywiper malware against Ukrainian local governments in 2022, aiming to destroy operational data [187]. Such incidents illustrate that municipal networks have become entangled in global geopolitical cyber conflict.

These developments reveal that municipalities face a uniquely complex cybersecurity challenge: they must protect citizen services, maintain public trust, and defend infrastructure increasingly targeted by both criminal and state-sponsored adversaries. It is therefore imperative that municipalities prioritize cybersecurity and evaluate security measures amidst this increasingly complex threat landscape.

## 1.1. BACKGROUND

To understand the challenges of implementing security measures within municipalities, we adopt the socio-technical causal model for cyber risk presented in [193, 274], as an analytical lens for examining cybersecurity. We further elaborate on this model in subsection 1.1.2. The socio-technical causal model offers a structured representation of the relationship between threat and harm. Municipalities themselves do not employ this analytical model for cybersecurity decision-making. Instead, they rely on security frameworks that provide actionable, structured guidelines and best practices for mitigating risks posed by threats and vulnerabilities. These frameworks, though practical and norm-setting, do not offer a systematic understanding of the underlying security challenges. Nevertheless, they currently serve as the foundation for most real-world municipal security efforts. Therefore, before introducing the causal model, we briefly outline several well-known frameworks.

### 1.1.1. SECURITY, FRAMEWORKS AND COMPLIANCE

For many municipalities, increasing security is equivalent to increasing compliance with security frameworks and standardized methodologies [163]. This is understandable, as municipalities need to balance limited resources and scarce expertise in the cyber domain, and security frameworks provide simple rules to follow, even without cyber expertise.

First, ISO/IEC 27001 provides an international standard for information security management. It helps organizations establish, implement, maintain, and continually improve an Information Security Management System (ISMS) and also adopts a risk-based

approach [82]. ISO/IEC 27001 specifies mandatory criteria that an organization, system, or process must meet to comply. Essentially, it provides a compliance framework that can be used as the basis for certification or an audit. Additionally, ISO/IEC 27002 provides a guidance standard that supports ISO/IEC 27001. It provides detailed best practices for implementing the security controls referenced in Annex A of ISO/IEC 27001. Whereas ISO/IEC 27001 is certifiable, ISO/IEC 27002 is not; it's purely a reference. It provides guidance on how organizations can implement security controls in practice, given their organizational context [83].

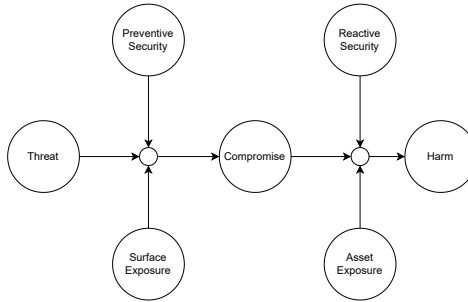
In the Netherlands, the standards ISO/IEC 27001 and ISO/IEC 27002 are the foundation of the national regulatory framework Baseline Informatiebeveiliging Overheid (BIO). This framework is used throughout all layers of the Dutch government to provide a uniform framework for information security. At the European level, the NIS2 directive aims to strengthen cybersecurity baselines across member states of the European Union (EU) [43]. It establishes a common level of cybersecurity risk management and incident reporting obligations for European critical sectors. For Dutch municipalities, the BIO serves as the foundation for information security management, with NIS2 compliance layered on top, adding reporting and accountability measures. However, while the BIO framework is in place to encourage security measures, there are few hard requirements for municipalities, and no enforcement is in place. As a result, the effectiveness of this framework remains unclear.

Germany has a similar approach to the Netherlands on regulatory frameworks for municipalities, with a longstanding baseline security standard from the Federal Office for Information Security (BSI). It provides catalogs of measures used by federal and public institutions, and is certifiable under ISO/IEC 27001 [96]. In the UK, multiple frameworks exist; however, the Cyber Essentials certification scheme is the minimal security baseline for any organization [31]. In the US, there is a federal baseline for federal agencies, NIST SP 800-53 (Security and Privacy Controls for Federal Information Systems), but it does not apply to municipalities. Instead, they are free to choose their own frameworks, including, among others, NIST CSF or CIS Controls, detailed below. Some municipalities use the NIST CSF [184], while others use the CIS Critical Security Controls framework [51].

As mentioned earlier, the BIO regulatory framework is specific to the Netherlands; internationally, other popular frameworks include the NIST Cybersecurity Framework [183] and the CIS Controls [81]. These frameworks, detailed below, are particularly adopted by municipal organizations in the US.

The NIST Cybersecurity Framework 2.0 (NIST CSF) is a voluntary framework that provides organizations with a risk-based approach to managing cybersecurity. It emphasizes that cybersecurity is not just an IT job; it is a risk-management issue for the entire organization. It does so by providing six “Functions” that serve as the framework's building blocks and represent the priorities of cybersecurity risk management. In providing these functions, it provides a strategic view of cybersecurity risk, elevating cybersecurity from a pure technical issue to the board for organizational preparedness and response mechanisms [183].

The CIS Controls are a prioritized, prescriptive set of safeguards designed to mitigate common cyber threats [81]. It is developed by the US non-profit Center for Inter-



**Figure 1.1:** The social-technical causal model as a structured way to understand municipal security challenges.

net Security (CIS), which develops cybersecurity best practices and guidelines. The CIS Controls offer a granular approach and are often practical for smaller organizations. CIS also manages the Multi-State Information Sharing and Analysis Center (MS-ISAC), which supports state and local US governments, including municipalities [80].

In sum, while these security frameworks provide hands-on, actionable, and norm-setting guidelines, successfully addressing cyber threats remains challenging in practice. Indeed, municipalities frequently fail to successfully deflect ransomware attacks, as evidenced by the number of incidents. Clearly, there is much to improve in the handling of municipal systems and processes to address security challenges. For example, before the ransomware attack on the city of Atlanta, the Atlanta government failed to upgrade its IT infrastructure, leaving it vulnerable to multiple security vulnerabilities. In January 2018, an audit found 1,500 to 2,000 vulnerabilities in the city’s systems [41]. Similarly, at Hof van Twente in the Netherlands, basic security measures were not in place and lacked monitoring capabilities, allowing criminals to dwell in the city’s systems for weeks [258]. These failings illustrate that security frameworks alone and mandatory norms are insufficient. Clearly, we lack empirical evidence on effective practices for improving security.

### 1.1.2. ANALYTICAL LENS FOR MUNICIPAL SECURITY: CAUSAL MODEL

To structurally understand security challenges for municipalities, we examine the socio-technical causal model [274], depicted in Figure 1.1. In this model, addressing cyber threats is understood as managing cyber risk [274]. While risk itself cannot be directly observed, its manifestation can be conceptualized and measured as “losses”, or more generally as “harm”. For harm to occur, a threat condition is required. Security can therefore be conceptualized as moderating the relationship between threat and harm, i.e., the interaction between an attacker and a victim organization. For a given threat level, more effective security should translate into lower expected harm. This high-level relation is represented by the leftmost and rightmost nodes in Figure 1.1.

The model further makes explicit that the relationship between security and harm is mediated by several intermediate factors. Reading the model from left to right, surface exposure captures the number and diversity of technical avenues through which a threat can attempt to compromise an organization, such as unpatched systems, misconfigurations, or externally accessible services. Greater surface exposure increases the likelihood

that a threat can successfully achieve a compromise, understood here as the violation of a security objective, such as unauthorized access, loss of integrity, or service disruption. Preventive security measures, such as patching, hardening, and access control, act at this stage by reducing surface exposure or increasing the effort required to achieve a compromise.

A compromise, however, does not automatically result in harm. The extent to which a compromise translates into harm depends, first, on asset exposure, which reflects the value and criticality of the affected assets. A compromise involving high-value assets poses greater potential harm than one involving assets of limited value. Second, reactive security measures, such as detection, containment, incident response, and recovery, shape how much harm ultimately materializes after a compromise has occurred. Effective reactive security can limit losses even when preventive measures fail. Together, these factors explain how threats, exposure, and security interact to determine expected harm in the socio-technical causal model [274].

This model offers a simple yet structured understanding that highlights the key areas shaping municipalities' security realities. Therefore, to adequately address cyber threats, municipalities should *i*) explore security measures (preventive and reactive) to minimize the effect of threats on their organization, *ii*) reduce their surface exposure to minimize the risk of compromise, and *iii*) know about the threats facing them. In the next sections, we iterate over these three key areas.

### 1.1.3. PREVENTIVE AND REACTIVE SECURITY MEASURES

Organizations employ a wide range of preventive and reactive security measures, such as technical controls, organizational practices, and human-centered interventions. This chapter does not aim to provide a comprehensive overview of all available security mechanisms. Instead, we focus on three measures that recur across the incidents discussed earlier in this chapter. For example, the incident in Atlanta showed that prior to the ransomware incident, an audit found roughly 1,500–2,000 vulnerabilities, highlighting an enormous backlog of delayed patches [55, 207]. Similarly, in the case of Hof van Twente [258], the attack occurred because basic security measures were not in place, indicating shortcomings in both preventive and reactive controls. Therefore, we examine the following three measures, *i*) patching behavior, *ii*) vulnerability notifications, and *iii*) incident response and Computer Security Incident Response (CSIRT) teams. The following sections briefly describe each focal measure and its specific role in municipal cybersecurity.

#### PATCHING AND VULNERABLE SYSTEMS

Patching is a preventive security measure that mitigates known vulnerabilities by keeping systems up to date. In many cyber attacks, exploiting known vulnerabilities for which a patch exists remains a dominant attack vector, even after years of warnings [37]. System administrators keep systems updated, and any computer system within an organization needs to be administered.

Because of their pivotal role in managing computers within an organization, it makes sense to understand the considerations of system administrators. System administrators operate within large-scale, complex environments that present significant technical, so-

cial, cognitive, and business challenges [20, 123, 264]. They often act as a broker between the end-users and the technical community [260]. Consequently, system administrators play a key role in managing organizations' computer systems. In the context of patching, system administrators face challenges in comprehensively acquiring meaningful information about available updates, effectively testing and deploying updates on time, recovering from update-induced problems, and navigating organizational and management influences [135].

In addition to the role of the system administrator, the location of a system within a network affects patching behavior. Patching server applications occurs much more slowly than patching client-side applications [120]. Closely related, vulnerability remediation is a struggle among many practitioners. Even discovering vulnerable systems is impeded by factors such as trust, communication, funding, and staffing [7]. Furthermore, vulnerability notifications are often so plentiful that 95% of CVE disclosures are not ingested by organizations [54].

Vulnerable systems that are not patched are quickly identified by (security) researchers and attackers alike. Nowadays, passive scanning services scan the Internet daily and provide a user-friendly interface to look up vulnerable systems. Examples of such services are Censys [58] and Shodan [150]. With such easy-to-use passive scanning tools available, it's easy to detect vulnerable hosts [24, 59, 63, 91]. However, existing methods to measure vulnerable systems are not without limitations. For example, the Internet-facing OpenSSH service might not be as vulnerable as initially suspected due to the use of backports [270].

#### VULNERABILITY NOTIFICATIONS

To know which systems to patch, an organization needs to learn about vulnerabilities. In doing so, a notification may be sent to an organization to alert them about the vulnerable system. Many different organizations may send vulnerability notifications, including vendors, sectoral or national Computer Security Incident Response (CSIRTs) teams, and independent researchers.

Notifications differ in their degree of specificity. Some are generic, meaning they do not identify whether a particular organization is affected. Vendor-issued communications about newly discovered vulnerabilities in a product fall into this category and are commonly referred to as advisories. In contrast, asset-specific vulnerability notifications indicate that a concrete system, such as a service reachable at a specific IP address, has been observed to be vulnerable. These notifications explicitly support targeted remediation by linking vulnerability information to an identifiable asset and are referred to as "vulnerability notifications". In the Netherlands, for example, sector CSIRTs send asset-specific vulnerability notifications to their constituents.

As a preventive security measure, sending vulnerability notifications intuitively makes sense. After all, preventing an incident is more effective than remediating one. Indeed, in practice, sending notifications increases patch rates at organizations [60, 134, 231, 252]. Sending notifications, though, does come with a variety of challenges. For example, obtaining contact details at scale for recipients is problematic [34]. Also, the content of a notification, rather than the sender's reputation, affects whether the recipient acts on it [280]. Finally, notifications sent directly to an organization appear to have a stronger

remediation effect than notifications sent to a national CERT, which forwards them [133]. It turns out that many exploitation attempts of critical vulnerabilities occur shortly after the announcement or disclosure of the vulnerability, suggesting that in the notification process, time is of the utmost importance [192].

#### CSIRTS AND INCIDENT RESPONSE

Computer Security Incident Response Teams (CSIRTS) help organizations prevent and mitigate cyber incidents. CSIRTS provide their services and support to a defined constituency. They manage information security incidents by preventing, handling (i.e., detecting, analyzing, responding), and/or coordinating information security incidents [77]. CSIRTS combine preventive and reactive security measures for a municipality. On the one hand, they help prevent incidents, while on the other hand, they assist when a municipality has been compromised.

The authoritative FIRST (the Forum of Incident Response and Security Teams) delineates various types of incident response teams, including CSIRTS, Product Security Incident Response Teams (PSIRTS), Security Operation Centers (SOCs), and Information Sharing and Analysis Centers (ISACs) [76, 77]. According to FIRST industry guidance, CSIRTS may be established as a single unit, an independent organization, or a part of a larger cybersecurity organization, like in many national cybersecurity centers (NCSCs). FIRST further notes that National CSIRTS and sectoral CSIRTS are special types of CSIRTS that coordinate responses to information security incidents, threats, and vulnerabilities. FIRST will describe sector CSIRTS in future versions of the framework [77].

Alongside FIRST, the CERT Division of Carnegie Mellon University developed a framework to establish a sector CSIRT [180]. Other industry efforts have largely focused on non-sectoral CSIRTS [13, 64–66, 230], guiding the establishment and operation of CSIRTS, as exemplified by the continuously updated handbook for CSIRT teams [112, 271].

Sector CSIRTS are on the rise in regulatory frameworks to tackle cyberattacks. In theory, sector CSIRTS perform the same functions as the national CSIRT, but then for a more targeted set of constituents. The constituents typically cover a specific sector such as finance, energy, or water. Institutionally, sector CSIRTS are a model that has been copied worldwide [180], and several studies have argued for the benefits of a CSIRT for a specific sector [97, 166, 267]. To channel information between various organizations, including the national CSIRT and constituents, sector CSIRTS maintain relationships between many stakeholders. These organizations have grown in importance over the last decade. In 2016, the US Presidential Policy Directive 41 (PPD-41) highlighted the importance of sector-specific cybersecurity measures and the collaboration between the Department of Homeland Security (DHS) and sector-specific CSIRTS [185]. The terminology varies across countries; in the US, sector CSIRTS are referred to as Information Sharing and Analysis Centers (ISACs). In the European Union, the Directive on Security of Network and Information Systems (NIS) mandated that certain sectors establish incident response capabilities, including sector-specific CSIRTS [67]. Its 2022 successor in the EU, NIS2, strengthens the role of sector CSIRTS by expanding the number of sectors required to establish a sector CSIRT from 7 to 18.

Trust and communication are critical components for CSIRTS. Trust, personal relationships, and information-sharing networks are essential for collaboration and successful incident management [13, 88, 124, 165, 272]. In the context of cybersecurity,

when trust is absent, information sharing of sensitive data is likely to be hindered [244, 269].

Human factors also play a significant role in the effectiveness of incident response teams. The success of these teams depends not only on technical capabilities but also on the dynamics of individuals working together [35, 200, 206, 228, 232, 255].

#### 1.1.4. SURFACE EXPOSURE: MANAGING MUNICIPAL ASSETS

The more systems an organization operates, the more potential attack vectors an attacker has. According to the causal model in Figure 1.1, each additional computer system increases the municipality's attack surface and thus the risk of harm, as reflected in "Surface Exposure". Minimizing the surface exposure of an organization is the area known as "asset management".

Today, organizations typically operate thousands of computer systems [191, 229], and obtaining an inventory of all those assets is a daunting task. Many technical solutions exist to automate the asset discovery process [72, 140, 218, 236, 262]. These technical solutions, however, are insufficient to capture the complexities that surround asset management. For example, asset management within organizations often fails due to policy compliance issues and so-called shadow IT: systems that are not in the official inventory of an organization [204]. Ad-hoc initiatives and circumventions, created by productivity-focused employees when the organization's existing security implementations do not meet their needs, may lead to shadow IT [113]. In practice, the key factors in the decision to comply with security rules are the actual and anticipated costs and benefits of compliance to the individual employee, and perceived costs and benefits to the organization [21]. Consequently, it's the multifarious roles of human and organizational factors, as opposed to solely technological programs or programming mistakes, that may lead to vulnerabilities and difficulties in asset inventory [121].

#### 1.1.5. THREAT: KNOW YOUR ATTACKER

In the causal model, "threat" represents entities attacking a municipality. It is a requirement for harm to occur [274]. This relation between threat and harm has found its way into security frameworks and best-practice guidance. Today, best practices propose that organizations focus their limited resources on tailored defences, i.e., knowing about the adversaries targeting them [48, 146, 182, 197, 224]. This, therefore, seems to be particularly useful for municipalities and their growing threat landscape, as it helps them prioritize security measures given their limited resources. Organizations may ingest threat intelligence (TI) to learn of adversaries' indicators of compromise (IOCs), and Techniques, Tools, and Procedures (TTPs). Known as 'threat-led defense', this movement found its way into government and corporate policy guidance [15, 247, 249].

Threat intelligence as a defensive measure is built on the assumption that sharing attacker infrastructure found at one organization may help protect another organization [107]. TI can be free or commercial. In both cases, while indeed promising, TI is not without limitations in accuracy, timeliness, and coverage [126, 152, 243]. Moreover, the quality of TI sources fluctuates when measured along various dimensions, including, among others, coverage, timeliness, relevance, overlap, delay, volume, and accuracy [25, 95, 118, 136, 220, 261]. In essence, it's really hard to measure quality without any

form of ground truth on attackers [152]. Yet, despite the availability of such metrics, they are often not relevant to TI consumers. The main purpose of TI for customers appears to be understanding the threat landscape, rather than detection [25]. Attributed threat intelligence, i.e., knowing exactly who is conducting attacks or operating malicious infrastructure, is very hard without ground truth [152]. Moreover, how exactly attribution is done by TI vendors is not well understood [216], but the reliability of attributed TTPs has been questionable [217, 254]. One difficulty in understanding this process has been the unreliable naming and labeling of threat actors by TI vendors, making attribution comparisons very difficult [93].

The causal model provides a structured framework for understanding the security challenges municipalities face. One strength of the model is that it does not focus solely on technical measures; it also captures organizational and institutional factors that shape municipalities' security realities. In the next section, the scientific state of the art is surveyed to identify research gaps in the articulated security measures. The following section specifies research aims and formulates the main research question and sub-questions.

## 1.2. RESEARCH GAPS

In section 1.1, we outlined the background on municipal security challenges and presented the causal model as a structured way to understand those challenges. We surveyed three key areas of security measures for municipalities to address cyber threats: *i*) preventive and reactive security measures, *ii*) managing surface exposure, and *iii*) knowledge about threats facing municipalities. In the next sections, we identify and describe three gaps in the scientific literature.

### PATCHING BEHAVIOR AND MANAGEMENT OF VULNERABLE HOSTS

Frameworks such as the regulatory BIO framework and the NIS2 legislation have emphasized the importance of patching. Yet, patching remains an unsolved problem for organizations in both the government and the private sector [37]. In particular, local governments appear to be struggling to develop adequate capabilities [102, 153]. For example, the U.S. Cybersecurity and Infrastructure Security Agency (CISA) referred to them as the “cyber poor”, offering vulnerability scanning services as support [132].

So, why do internet-facing systems go unpatched if the importance of this issue is clear? Prior work has left a gap in how organizations, particularly municipalities, manage vulnerable systems. For example, previous work showed that vulnerable hosts can be observed from outside an organization via passive scanning services [24, 59, 61, 63, 91]. However, external measurements do come with limitations [270]. Second, earlier work shed light on the role of system administrators in the patch process [20, 56, 260, 264]. In [135], the authors identified patch processes, however, the software updates that admins reported on in this study, though, lacked empirical validation. These studies examined the perspectives of relevant practitioners but relied on self-reported behaviors, potentially biasing the results. Moreover, existing studies have failed to address the challenges that practitioners in municipalities [191, 229] face in managing municipal assets.

Prior work, it seems, has not connected vulnerability scanning research with user

studies of system administrators. Perhaps the detection of vulnerable systems, i.e., unpatched systems, is inaccurate. Or, system administrators might be unaware of the systems. Or, a deliberate decision was made not to patch.

#### INSTITUTIONAL SUPPORT STRUCTURES FOR INCIDENT PREVENTION AND MITIGATION

Industry guidance [76, 77, 180] and academic studies on CSIRTs have largely focused on non-sectoral CSIRTs [13, 64–66, 230]. At the same time, the need for sector CSIRTs across sectors, including municipalities, is growing [97, 166, 267]. And while the role of a sector-based approach in regulatory frameworks is growing [67, 185], our understanding of these organizations and the challenges facing practitioners operating them is not. In the context of vulnerability notifications, sector CSIRTs play a pivotal role in the Netherlands, as they distribute the national CSIRT's notifications to sector-specific constituents. No work has yet addressed the challenges of operating a sector CSIRT for municipalities or other sectors, or examined how the vulnerability notification mechanism involving CSIRTs, i.e., the national CSIRT and a sector CSIRT, is functioning.

#### QUALITY EVALUATION OF ATTRIBUTION IN COMMERCIAL TI

Frameworks, policy guidance, and best practices advocate that organizations know the adversaries that are targeting them [15, 48, 146, 182, 197, 224, 247, 249]. In doing so, organizations need to rely on attributed TI. While attribution seems to take center stage in policy, not much is known about the coverage and reliability of attribution by TI vendors. Disagreement among vendors may undermine the actor-centric defenses that are proposed. There is an abundance of earlier work on TI [95, 118, 126, 136, 152, 220, 243, 261], but it rarely focuses on attribution. Most of the previous work has focused on open TI, i.e., freely available TI. Only one study has looked at commercial (paid) TI [25]. However, to date, no study has systematically investigated attribution for IOCs of a large set of commercial TI vendors.

In sum, we identify the following three gaps in state-of-the-art scientific literature:

1. We lack insights into why known vulnerable internet-facing systems remain online.
2. We do not have an understanding of how institutional support structures for incident prevention and mitigation are functioning.
3. There is no independent evaluation of the quality of attribution in commercial threat intelligence.

We conduct three research activities to fill these research gaps. First, we measure vulnerable hosts in Dutch municipalities via passive scanning services. With an established set of vulnerable hosts, we then interview practitioners responsible for managing them to determine the root causes of those systems not being patched. For our second research gap, we take a look at the “Informatiebeveiligingsdienst” (IBD): the Dutch sector CSIRT responsible for municipalities. We use the IBD as a case study to determine the expectations and challenges in providing sector CSIRT services. We identify three stakeholder groups relevant to a sector CSIRT. We interview participants from each stakeholder group to measure their expectations on CSIRT services and the challenges in

providing them. Finally, we resolve the third research gap by conducting a longitudinal comparative analysis of unique, attributed IOCs collected from seven commercial TI feeds. In doing so, to evaluate attribution quality, we measure the level of agreement among the IOCs attributed by two or more TI vendors.

### 1.3. RESEARCH AIMS AND QUESTIONS

This dissertation examines the security measures municipalities use to address cyber threats. It investigates how municipal security measures are functioning and, using those insights, how those measures can be improved.

*How can municipalities improve security measures to address cyber threats?*

The focus of this dissertation is on the functioning of security measures that municipalities can employ to address cyber threats. Beyond simple security guidelines in existing security frameworks, we seek to understand the relationships among technical, human, organizational, and institutional factors that affect municipal security. We leverage technical work produced by other scholars and combine technical measurements with practitioner perspectives to develop a socio-technical perspective that better engages these security measures.

The scientific contribution of this dissertation lies in the combination of empirical results from the mixed-methods used to gain novel perspectives on the functioning of municipal security measures and the derived actionable insights. These perspectives aid practitioners in effectively implementing security measures and policymakers in formulating policies to address cyber threats against municipalities.

The three studies that make up this dissertation are detailed in the next section. Given our technical and human-centred approach, our research demands methods and analyses that, at heart, are multi-disciplinary. Therefore, we employ both quantitative and qualitative methods throughout the studies. Given that this is a paper-based dissertation, the study-specific research methodologies will be elaborated on in the individual chapters.

### 1.4. DISSERTATION OUTLINE

This section provides an outline of the dissertation. It describes three studies, their corresponding research questions, and presents an overview of the peer-reviewed papers associated with each study and chapter. Table 1.1 lists this overview, including the names of the collaborating researchers.

#### STUDY 1 - FACTORS IMPACTING PATCHING BEHAVIORS AND ASSET MANAGEMENT

Many organizations continue to expose vulnerable systems that have available patches, leaving them open to cyberattacks. Local governments are found to be especially affected by this problem. Why are these systems not patched? Prior work relied on vulnerability scanning to identify unpatched systems, notification studies on remediating them, and user studies of sysadmins to describe self-reported patching behavior, but these are rarely used together as we do in this study. We analyze scan data following

standard industry practices and detect unpatched hosts across the set of 322 Dutch municipalities. Our first question is: Are these detections false positives? We engage with 29 security professionals working for 54 municipalities to collect ground truth.

All detections were accurate. Our approach also uncovers a major misalignment between the systems that the responsible CERT attributes to municipalities and the systems that practitioners in municipalities believe they are responsible for. We then interviewed the professionals to find out why these vulnerable systems were still exposed. We identify four explanations for non-patching: unaware, unable, retired, and shut down. The institutional framework to mitigate cyber threats assumes that vulnerable systems are first correctly identified, then correctly attributed and notified, and finally correctly mitigated. Our findings show that the first assumption is correct, the second is not, and the third is more complicated in practice. We end with reflections on how to better remediate vulnerable hosts.

The research questions addressed by this study are as follows:

**RQ1:** How accurate are the measurements of unpatched systems?

**RQ2:** Why are those systems not patched?

#### STUDY 2 - FUNCTIONING OF INSTITUTIONAL SUPPORT STRUCTURES FOR MUNICIPAL INCIDENT PREVENTION AND MITIGATION.

In this paper, we study the experiences of practitioners in sectoral Computer Security Incident Response Teams (CSIRTs)—specialized teams that mediate between national cybersecurity authorities and the sector constituency. Through interviews with 18 professionals connected to the Informatiebeveiligingsdienst (IBD-CSIRT) for Dutch local governments, we uncover tensions in how key services are valued. For vulnerability notifications, while CSIRT staff consider them a core service, many constituents hardly mention them, and systemic gaps in information forwarding mean that crucial alerts often never reach them. We extend these insights with 5 interviews across other sector CSIRTs and a validation workshop with 7 participants, all security officers from sector CSIRTs, revealing shared challenges in balancing technical expertise with sector knowledge, building trust-based relationships, and navigating institutional bottlenecks. Our findings contribute the first systematic account of how sector CSIRT professionals understand and perform their role, highlighting the tensions in providing sector-wide support to professionals with differing security needs.

The research questions addressed by this study are as follows:

**RQ1:** What are service-specific challenges and expectations of stakeholders (sector CSIRT staff, governance bodies, and constituents) on the services provided by CSIRT practitioners?

**RQ2:** What are the strategic challenges for practitioners of a sector CSIRT in providing these services?

### STUDY 3 - QUALITY EVALUATION OF ATTRIBUTION IN COMMERCIAL TI

Attributed cyber threat intelligence (TI) plays an important role in mitigating cyber attacks effectively. Yet, despite the central role of attribution in policy, practice, and vendor reporting, little is known about the coverage and reliability of threat intelligence vendors' attribution. No study has systematically investigated attribution across a large set of leading TI vendors. We close this gap and provide a longitudinal comparative analysis across 13.5 million IOCs collected over the last 14 years from seven vendors. To compare IOC attribution across vendors, we normalize heterogeneous feeds and reconcile actor names using an evaluated and augmented version of MISP Threat Actor Galaxy (MISP TAG). Next, we address two questions: (i) what is the scope of actor-tracking by vendors, and (ii) how consistent is attribution among vendors? We find that the majority of actors tracked by one vendor are not tracked by the other. Furthermore, IOCs observed by multiple TI vendors are rare (1%), illustrating that commercial TI feeds, like open-source feeds, primarily provide singleton IOCs. We also find limited overlap in IOCs for jointly tracked actors by two vendors. We measure attribution agreement among vendors with Krippendorff's  $\alpha$ . We find poor agreement among vendors for actor attribution. By contrast, country attribution has high agreement. When removing 'temporary' attributions, agreement in attribution increases. Our results have implications for actor-centric defenses, compliance, and geopolitical uses of attribution.

The research questions addressed by this study are as follows:

**RQ1:** What is the scope of actor-tracking by vendors?

**RQ2:** How much agreement is there in attribution across vendors?

To reiterate, Table 1.1 lists this overview of the chapters in the remainder of this dissertation and the peer-reviewed studies they are based on. In Chapter 5, this dissertation is completed with a summary of the main findings, a reflection on the results, and future research directions.

**Table 1.1:** Dissertation outline.

Chapter	Publication
Ch. 2	<b>Ethembaoglu, A.M.</b> , van Wegberg, R.S. & Zhauniarovich, Y. & van Eeten, M.J.G. (2024). “The Unpatchables: Why Municipalities Persist in Running Vulnerable Hosts”. In <i>Proceedings of the 33rd USENIX Security Symposium (USENIX Security '24)</i> . [70]
Ch. 3	<b>Ethembaoglu, A.M.</b> , Kadenko, N.I., & Angelova, Y., & Zhauniarovich, Y. & Parkin, S. & van Eeten, M.J.G. (2026). ““Tell Them They Are a Responsible Entity, Not a Customer”: Understanding Practitioner Challenges in Sector CSIRTs”. In <i>Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)</i> [69]
Ch. 4	<b>Ethembaoglu, A.M.</b> , van Wegberg, R.S. & Zhauniarovich, Y. & van Eeten, M.J.G. (2026). “APT to Disagree: A Comparative Analysis of Attribution in Commercial TI”. In <i>Proceedings of the IEEE Symposium on Security and Privacy (S&amp;P '26)</i> .

# 2

## THE UNPATCHABLES: WHY MUNICIPALITIES PERSIST IN RUNNING VULNERABLE HOSTS

*Many organizations continue to expose vulnerable systems for which patches exist, opening themselves up for cyberattacks. Local governments are found to be especially affected by this problem. Why are these systems not patched? Prior work relied on vulnerability scanning to observe unpatched systems, notification studies on remediating them, and on user studies of sysadmins to describe self-reported patching behavior, but they are rarely used together as we do in this study. We analyze scan data following standard industry practices and detect unpatched hosts across the set of 322 Dutch municipalities. Our first question is: Are these detections false positives? We engage with 29 security professionals working for 54 municipalities to collect ground truth.*

*All detections were accurate. Our approach also uncovers a major misalignment between systems that the responsible CERT attributes to the municipalities and the systems the practitioners at municipalities believe they are responsible for. We then interviewed the professionals as to why these vulnerable systems were still exposed. We identify four explanations for non-patching: unaware, unable, retired and shut down. The institutional framework to mitigate cyber threats assumes that vulnerable systems are first correctly identified, then correctly attributed and notified, and finally correctly mitigated. Our findings illustrate that the first assumption is correct, the second one is not and the third one is more complicated in practice. We end with reflections on how to better remediate vulnerable hosts.*

---

This chapter has been published as: **Ethembabaoglu, A.M.**, van Wegberg, R.S. & Zhauniarovich, Y. & van Eeten, M.J.G. (2024). "The Unpatchables: Why Municipalities Persist in Running Vulnerable Hosts". In *Proceedings of the 33rd USENIX Security Symposium (USENIX Security '24)*.

## 2.1. INTRODUCTION

Exploiting known vulnerabilities for which a patch exists remains a dominant attack vector, even after years of warnings [37]. Local governments are seen as especially susceptible [102, 153]. The U.S. Cybersecurity and Infrastructure Security Agency (CISA) referred to them as the “cyber poor”, offering vulnerability scanning services as support [132]. These concerns are not exclusive to the U.S. In the Netherlands, the Dutch Safety Board investigated the incidents following the 2020 Citrix vulnerabilities and concluded that municipalities struggle with patching because of a lack of resources [62].

The threat of exploitation of local governments, or any other organization, is not hypothetical. Municipalities worldwide have been hit with paralyzing ransomware attacks [36, 94, 179, 196]. These attacks had destabilizing societal effects, with governmental services being unavailable and data of citizens being lost. In the US alone, more than a hundred local government organizations reported cyberattacks in 2019 and 2020 [176].

To mitigate such threats, governments established Computer Emergency Response Teams (CERTs) or Computer Security Incident Response Teams (CSIRTs) [74]. These organizations receive security incident data and network scan information from various sources. This data is forwarded to the organization responsible for the vulnerable systems. The notified organization is then expected to mitigate the vulnerability. Prior research shows that such security notifications can expedite vulnerability remediation [133, 252]. CERTs around the world operate on a similar model of monitoring networks and notifying constituents. In Brazil, the CERT provides incident analysis and coordination services for any network that uses IP addresses or Autonomous Systems allocated to Brazil, and domains under the .br ccTLD. It alerts Brazilian networks involved in malicious activities [33]. The CERT-Bund in Germany supports handling IT security incidents; it provides active alerts for the federal administration in the event of acute threats [79]. In Africa, the non-profit organization AfricaCERT, with several African countries as its members, states in its objectives that it “encourages information sharing in ICT security, which includes findings from reported incidents and case studies, so that vulnerabilities can be rapidly identified, and its risks mitigated” [3]. In the UK, the NCSCS introduced the Early Warning service, offering its members to notify them of vulnerabilities in their networks [170]. These CERTs monitor threats and attribute IP networks to organizations for alerting and notifications.

Yet, despite these measures, governmental systems, as well as other organizations, are still frequently found to be vulnerable [238]. Therefore, in this study, we first ask the question: how accurate are the measurements of unpatched systems? Next, we consider the question: why are those systems not patched? Prior work has shed light on the presence of unpatched systems via three strands of research, (*i*) studying the self-reported patching practices of system administrators [56, 135], (*ii*) network scans that detected where vulnerable systems are located [120, 270], and (*iii*) the effectiveness of notifications about detected vulnerable systems [133, 252]. However, no prior work has integrated these three strands as we do in this study.

We combine passive network scans (relying on banner-grabbing to infer software versions) to link versioned software to known vulnerabilities from the National Vulnerability Database (NVD) and to detect vulnerable systems in local governments in the Netherlands that receive notifications from their CERT and then use these detections in

interviews with security practitioners responsible for those systems. We gather ground truth on the detections to answer our first research question. Next, to answer our second question, we explore the non-patching behavior of practitioners in a way less sensitive to the biases that come with self-reporting, which previous studies have relied on, e.g., [135].

We collaborated with IBD-CERT, the CERT organization for all municipalities in the Netherlands. The municipalities have registered their IP network ranges with the IBD-CERT. The CERT receives scan data from the national CERT and other sources about hosts with CVEs (Common Vulnerabilities and Exposures). It then notifies its members about detected hosts in their networks. We explore potential explanations for the detected vulnerable systems. First, the passive network scans might produce false positives: the host might not actually be running the vulnerable software. Scan data can contain artefacts and version information from hosts that are manipulated or simply wrong. Second, the systems are unpatched, but there is a reason the municipality has not patched it. It might be unaware of the vulnerability, it might be unable to patch it, or it might have decided that patching is not needed. In the process of conducting the interviews, a third explanation arose: the municipalities do not consider the vulnerable system to be their responsibility, even though they reside in the IP ranges that they registered with the CERT.

We analyzed 1,687 registered IP ranges covering 322 municipalities in the country. Using passive scanning data from Shodan [150] and Censys [58], we observed 154 vulnerable hosts running 17 different services with 643 unique CVEs in total. We conducted 16 semi-structured interviews with 29 security practitioners working for 54 municipalities (some IT departments support multiple municipalities), covering about 17% of the total population of Dutch municipalities. This sample includes municipalities with and without exposed vulnerabilities, so we can compare their answers and the features of their organizations. We transcribed, coded, and analyzed the interview data and conducted follow-up conversations.

First, we observed that the observation of vulnerable hosts seems reliable and not plagued by false positives. Next, we found that a significant portion of the vulnerable systems that get notified about fall into a gap, because there is a misalignment between the IPs of municipalities registered at the CERT, and the IPs the sysadmins see themselves as responsible for. At least some of these vulnerable systems appear to be “shadow IT”. This might explain why many of these systems persist in a vulnerable state. It also explains that the municipalities see themselves as much less vulnerable than their CERT or central government does. For the systems that were administered by the municipality, we observed that respondents were *(i)* not aware of vulnerable hosts, *(ii)* unable to patch the system, or *(iii)* systems were in the process of being retired. We also learned that vulnerable systems are rarely shut down because of security reasons. We make the following contributions:

- We collect ground-truth evidence that the external detection of vulnerable services is accurate and not plagued by false positives. The observations we collected from Shodan and Censys appeared to be 100% correct.
- We demonstrate major misalignments between the systems the CERT attributes

to a municipality and the systems a municipality believes it is responsible for. For our sample of municipalities, this misalignment translates to the CERT observing 18 vulnerable hosts that the IT departments do not consider their responsibility, pointing to the problem of “shadow IT”. On the other hand, the municipal IT departments do see themselves as responsible for 6 vulnerable hosts that the CERT does not attribute to them and thus doesn’t notify them about. Only 9 vulnerable hosts are seen and attributed consistently by both organizations. These observations raise concerns about the effectiveness of the incident response framework of CERTs for notifying victim organizations.

- We identify four categories for not patching vulnerable systems in practice: *unaware*, *unable*, *retired*, and *shutdown*. In most cases, security professionals were unaware of the vulnerable system. Additionally, we find that there are no CVE or application-specific mitigation strategies applied to vulnerable hosts unless explicitly provided in the security advisory of the CERT or from the vendor. We also observe a strong tension between business continuity and security in the vulnerability management process.

## 2.2. RELATED WORK

Our study ties into three main strands of research: (i) user studies of security practitioners or IT professionals; (ii) studies using network scans to collect observational data on vulnerable systems; and (iii) studies on security and vulnerability notifications. We discuss each in turn.

First, several studies examined the perspectives of security practitioners to gain an understanding of their considerations or the organizational processes in which they operate, not necessarily related to security [20, 264]. Li et al. [135] identified processes system administrators use to manage software updates but relied solely on self-reporting via surveys. The software updates that system administrators reported were not empirically measured, and therefore a picture may be painted that does not fully align with reality. Dietrich et al. [56] looked at how system administrators managed their systems and their configurations, specifically examining the perspectives of system administrators. Velasquez et al. [260] looked at the role of the system administrator within the organization and found that they often act as a broker between the end-users and the technical community. Kromholz et al. [123] found that the deployment process of security measures for system administrators is too complex and recommended that server configurations should opt for security by default. Alomar et al. [7] observed that practitioners struggle with vulnerability remediation and that vulnerability discovery efforts are hindered by significant trust, communication, funding, and staffing issues. Smale et al. [54] found that vulnerability information acquisition by practitioners is not comprehensive and that up to 95% of all CVE disclosures are not ingested. These studies examine the perspectives of practitioners, but they are all based on self-reported behaviors, which is potentially biased. By using actual network data linked to known vulnerabilities, we ground the responses of the interviewees by discussing with them the evidence of vulnerable hosts in their network.

Vulnerabilities can be found in the wild with passive scanning services. The work

of O'Hare et al. [186] presents a method to discover vulnerabilities by combining the CPE from passive scanning services with data from the National Vulnerability Database (NVD). A vast area of research relates to the use of Internet-wide scans with Zmap [61] and similar tools to detect vulnerable hosts [63, 91]. Numerous studies used Censys and Shodan to measure vulnerable systems [24, 59]. The work by West et al. [270] found that the Internet-facing OpenSSH service might not be as vulnerable as initially suspected due to the use of backports. Kotzias et al. [120] observed that the patching of server applications is much slower than the patching of client-side applications.

Finally, as stated in the introduction, there exists a line of research related to incident response and victim notification [252]. Li et al. observed that vulnerability notifications addressed directly to the owners of the resources promoted faster remediation than those sent to national CERTs [133]. Cetin et al. showed that retrieving contact information at scale was problematic. But once contacted, entities were more likely to remediate [34].

Our work builds on previous studies by connecting and contextualizing different methods and data sources to provide a deeper understanding of patching behavior and the responsibilities of networks. The combination of scanning networks and using that data in interviews should mitigate the risk of self-reporting in patching behavior. Additionally, it allows us to verify external measurements of software versions of networks, to obtain ground-truth. We complement those external network measurements with qualitative data to record the considerations of practitioners as to why those vulnerable systems exist in their infrastructure. Lastly, by collaborating with the CERT and the municipalities we are able to correlate (assumed) responsibilities for specific IP addresses that are attributed the municipalities.

## 2.3. ETHICS

We received approval from our Institutional Review Board for conducting this human-subjects research. Participants were explained in detail about the study, associated risks, and use of information for which they provided informed consent.

Research on the vulnerabilities of an organization is a sensitive topic. In our informed consent form, as well as in the recruitment emails for the interviews, we consistently assured the participating municipalities, as well as their CERT, that their data was handled confidentially and would only be presented in an aggregated and anonymized form. No identities or municipalities would be named. We also made sure that answers that referred to specific tooling that might reveal their identity were cleaned.

After the interviews, we reported IPs with vulnerable services to the CERT. During the interviews, we notified respondents of the observed vulnerable hosts.

To minimize the burden on the security staff, we choose to use passive rather than active scans in order to prevent the disruption of regular operations or unintentionally triggering alerts (false positives) in their Security Operation Centers (SOCs). Before publication, we presented a draft of this paper to the respondents, so that they had a chance to check and correct quotes attributed to them.

## 2.4. MEASUREMENT APPROACH

We used a mixed-methods approach that combines external network measurements of the municipalities with a qualitative user study among sysadmins and security practitioners. The aim of this approach is two-fold. First, we validate the passive network measurements for detecting software versions with the responsible operators. This allows us to estimate the false positive ratio for the detection of vulnerable systems when relying on banner information about software versions and linking those to known vulnerabilities. We did not carry out active measurements on the networks of municipalities to verify vulnerabilities. Second, we conducted interviews to learn from the practitioners why the vulnerable systems, assuming they were correctly detected, were present in their network. We collaborated with the IBD-CERT, which is responsible for all Dutch municipalities. The CERT provided us with the IP ranges that the municipalities have registered with it. It also facilitated the process of recruiting interviewees.

### 2.4.1. SCANNING MUNICIPAL NETWORKS

Network scans can be done actively or passively. Active scans directly connect to the target network. They grab banners and infer software versions (e.g., Nmap [139]) but they can also be more intrusive. Some tools such as Metasploit [201], Nessus [240] or Qualys [198] actually try to exploit a potential vulnerability to determine if it is present on the target system.

By contrast, passive scanners run their own Internet-wide banner-grabbing scans and present the results to its users, often as a service via a web portal. The scan data is (somewhat) outdated but the target network is not directly touched by the users of the service. Popular passive scanning services are Shodan [150] and Censys [58].

The CERT provided the research team with 1,687 IP ranges for 322 municipalities. In the Netherlands, there are 346 municipalities, giving us coverage of over 93% of the total population. To reduce the burden on the municipality networks, we relied solely on passive rather than active scans. In November 2022, we queried Censys and Shodan to identify hosts. This process resulted in 3402 detected hosts. Figure 2.1 depicts the 322 municipalities from the CERT and the number of responding hosts for each. The next step was to determine what services were running on the detected hosts. Censys and Shodan parse banners to determine the service and version that is running on a port. A portion of the service banners contained version information. The service and, if available, version are used to generate a Common Platform Enumeration (CPE) identifier. We collected CPE identifiers of the applications and their versions running on a host. For those CPEs, we retrieved the accompanying Common Vulnerability Disclosures (CVEs) from the National Vulnerability Database (NVD) [174]. For these CVEs, we also collected their Common Vulnerability Scoring System Version 3 (CVSS3) scores. CVSS3 is an open framework for communicating the characteristics and severity of software vulnerabilities [75]. We used the CVSS score to label vulnerabilities as Critical, High, and Medium/Low [175].

Our study primarily relies on the scan results from Censys. Contrary to Shodan, it does not perform any black-box post-processing, so we can more transparently infer CPEs. Furthermore, Censys runs its scanners on a daily basis for each IP address [29]. To corroborate its results, we compared the results to those we got from Shodan. We

observed two minor discrepancies. First, we found that for some IP ranges, each service returned a different number of hosts. (Both services did return the same number of hosts that were running a versioned service.) We asked respondents if they actively blocked either Censys or Shodan, which none did. However, several respondents did mention that their firewall blocks consecutive requests from a scanner and that it might explain the difference in resulting hosts. Second, we found minor discrepancies in how Shodan and Censys parse the banner for the CPE. For example, Shodan detects unversioned Apache2 instances conveying the following CPE for them: *cpe:2.3:a:apache:http\_server:2*, while Censys returned the *cpe:2.3:a:apache:http\_server:\*.\*.\*.\*.\*.\** CPE. Also, for Nginx 1.18.0, Shodan returns the CPE *cpe:2.3:a:igor\_sysoev:nginx:1.18.0* while Censys returns *cpe:2.3:a:nginx:nginx:1.18.0:\*.\*.\*.\*.\*.\**. For neither CPE did it have an effect on the associated vulnerabilities.

### 2.4.2. USER STUDY

**Selection of municipalities.** We compiled a list of versioned and (vulnerable) services per municipality. For the interviews, we selected municipalities on three criteria. First, we preferred municipalities with the highest number of hosts running versioned services, to maximize the number of external measurements of systems. We interviewed 10 municipalities with vulnerable hosts. Second, we interviewed municipalities without

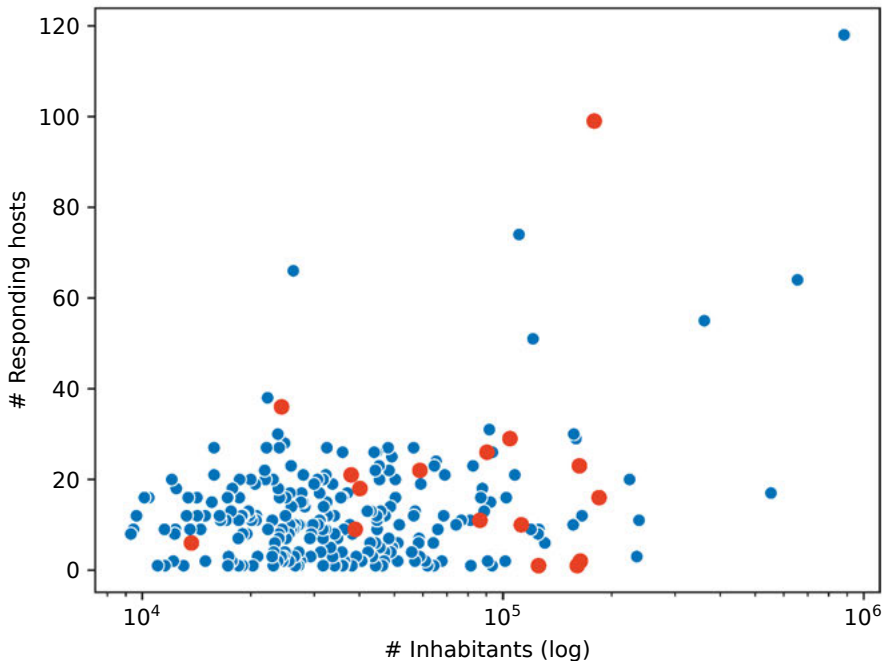


Figure 2.1: Responding hosts versus the number of inhabitants (log scale) for the IPs of all municipalities. Red dots are the municipal IT organizations (16) we interviewed.

vulnerable systems. We wanted to learn if they approached vulnerability management differently, or if their organizational structure may have an effect. We interviewed 6 municipalities without vulnerable hosts. Third, we wanted a diverse set of municipalities in terms of size and geographical location. We aimed for a diverse sample measured in the number of inhabitants and Internet-facing hosts. Depicted in Figure 2.1, we plotted in red the participating municipalities among the total set of municipalities.

Based on these criteria, we selected 34 municipalities and, in collaboration with the CERT, reached out to them in November 2022 until April 2023 via e-mail, inviting them to join a one-hour semi-structured interview. In case of no reply, we sent out a reminder after several weeks. We asked for respondents who were involved in the operational process of detecting and mitigating vulnerabilities, often system administrators. In total, we were able to set up interviews (4 declined to participate, and 12 did not respond). In several cases, municipalities maintained a shared ICT infrastructure with neighboring municipalities. In total, we interviewed 29 practitioners occupying 4 kinds of roles (Table 2.1) belonging to 16 organizations providing IT services to 54 municipalities (about 17% of the total population of Dutch municipalities). See Appendix A.1 for more information on respondents.

**Table 2.1: Respondent roles**

Roles	Respondents ( $n = 29$ )
System Administrator	9
Network Administrator	2
Security Engineer/Officer	12
(C)ISO	6

**Pre-interview engagement.** Our interviews were one part of a multi-step engagement with the responding municipalities. We first reached out to municipalities via the CERT contact points. We explained the purpose and design of the study and asked them to participate in an interview. The interview had two main purposes: to verify the validity of the detected vulnerable hosts and to understand the reasons for those hosts being present in the network. Discussing and validating specific hosts meant we needed to enable the municipalities to prepare for the interviews. First, we updated our scan results shortly before the engagement. Then we sent a list of hosts, services, and versions that we had detected to our main contact point at each municipality. The contact points then had to identify within their organizations which IT practitioners were responsible for the specific hosts that we wanted to discuss. The practitioners, in turn, would then be able to prepare for the interview, e.g., by checking the exact versions of the services running on the hosts.

**Adapting the Research Design.** A surprise emerged during the first few interview preparations. Some respondents told us that the hosts we had sent them were not known to them or were not their responsibility, even though they were located in the IP ranges that the municipality had registered with the CERT. They are responsible for updating the data about the corresponding IP ranges if changes happen, yet we consistently ran

into discrepancies. We see these discrepancies as outcomes of our research and will discuss this issue in depth in Section 2.5. That being said, it also meant we had to adapt our research design.

Our first adaptation was to ask respondents to share with us the IP ranges that they were responsible for. This brought another surprise: some of these ranges were completely outside of the ranges that were registered with the CERT. This led to a second adaptation: we asked all respondents to share with us, well before the interviews, the IP ranges they were responsible for. We would then scan these ranges, in addition to the ranges that the CERT had on record, for hosts and services in Censys and Shodan. We would also identify vulnerable hosts in these ranges. In this way, we could share up-to-date scan results in preparation for the interview, as well as ensure that we could discuss vulnerable systems. These additional scans brought into focus potentially vulnerable systems that were not attributed by the CERT to the municipality.

**Interview protocol.** The combination of conducting interviews about actual scanned vulnerable hosts provides an empirical basis for determining how practitioners manage vulnerable systems and reduces the risk of social-desirability bias that may occur with self-reporting. Previous works [56, 135] used qualitative data but did not relate that data to actual systems.

The interviews consisted of three parts. First, we would ask the interviewee to confirm or reject the version information we had inferred about the selected hosts and services. In other words, it acted as a ground-truth validation protocol for the vulnerable services and determined if the external measurements were indeed correct. Second, it seeks to understand practitioner perspectives on the responsibility of IP addresses for monitoring and mitigating vulnerable systems of an organization. Third, it determines how practitioners assess and mitigate vulnerable systems in practice. These discussions centered around the vulnerable systems that we measured, not hypothetical cases. For the latter part, we chose a semi-structured protocol because we wanted the interviewees to freely express their thoughts on the reasons for the presence of vulnerable hosts [98, 129]. The full interview protocol is included in Appendix A.2.

The interviews were conducted in person or using video conferencing applications and typically took about an hour. There were 6 interviews done with a single respondent, 5 with two, and 4 with three respondents. In those cases, respondents stated that several people were involved in administering the infrastructure and to improve our understanding, all should provide input. One municipality answered interview questions via e-mail because the infrastructure administration required too many people for one interview. In the pre-interview communication, the respondents were informed about the goals of the interview and received an Informed Consent form. At the start of the interview, we reiterated the research goals, gathered consent statements, and asked permission to record the interview for transcription. Respondents participated voluntarily and did not receive compensation for the interview.

**Coding.** Interviews were transcribed and coded using the ATLAS.ti software [12]. Initial codes were iteratively developed by the lead researcher and two other researchers. The codes clustered topics that described the various kinds of answers of participants. First,

4 interviews were coded by the lead researcher for the initial codes. The research team then refined the initial codes. These codes were then shared with another researcher to independently code a subset of interviews and discuss the results. We refined the codes with the research team with those results, leading to the final codebook. The process of meeting with authors and discussing and independently refining the codebook is a suitable way to ensure the reliability of findings, according to McDonald et al. [151]. The final codebook is available in Appendix A.3.

## 2.5. VALIDATING OBSERVED VULNERABLE SYSTEMS

The first potential explanation for the presence of vulnerable hosts is that their detection might contain false positives. That is, the network scan data received and disseminated by CERTs might not be fully accurate. Our approach, extracting CPEs from version information in banners, is a normal industry practice, so our data is similar to the data CERTs receive. There is a second type of scanning that does not rely on version information alone and instead uses benign exploit code to test the presence of a vulnerability. Because of the intrusive nature of such scans, we did not adopt this approach. Clearly, the second approach offers greater reliability for estimating vulnerabilities but the approach has a direct impact on the target network, which was unfeasible for our collaborations.

There are two caveats to the approach we used. First, administrators may hide the version in the banners they expose. Second, a service may run a backport, i.e., an older version that includes security patches from a new version but still shows the old banner. A banner rarely shows the presence of a backport of a service. Therefore we asked respondents if they used them. None of our 29 respondents indicated that they (knowingly) ran backports. In two interviews, respondents stated they use a security product that hides version information.

Our scans resulted in 3,402 records from Censys, i.e., that is the number of responding hosts. Within this dataset, 578 hosts were found with at least one versioned service application (17%). A host can run multiple (vulnerable) services on different ports. We, therefore, examine all unique vulnerable CPEs on all hosts. From the 578 hosts, we derived 101 unique CPEs. Of these 101 unique services, 70 contained 1 or more CVEs, with a total of 643 unique CVEs. The 70 vulnerable services were observed in 154 unique hosts in 94 different municipalities. In our population, vulnerable hosts most frequently ran vulnerable versions of OpenSSH and Apache Httpd.

To verify the services and versions at vulnerable hosts, we prepared a list of detected systems for each of the 54 municipalities, based on the IP ranges registered with the CERT. In total, this set contained 24 vulnerable hosts for the 16 IT organizations servicing 54 municipalities.

However, as explained in Section 2.4.2, the pre-interview communication revealed that the municipalities did not consider some of the CERT-registered IP ranges as falling under their responsibility. So some of the detected vulnerable systems were not under their control or even unknown to our interviewees. This meant that in 9 of the 16 interviews, there were no vulnerable hosts at the municipality in the IP ranges reported by the practitioners. This misalignment is explored further in Section 2.6. In the remaining 7 interviews, we were able to validate 15 vulnerable hosts. We supplemented this set by validating the 27 non-vulnerable hosts since that inference (vulnerable or not) is based

on the exact same data and analysis. This allowed us to test whether using version information from banners is reliable or plagued by a substantial false positive rate. The 15 vulnerable hosts ran 4 different vulnerable services: Apache, OpenSSH, PowerDNS, and Nginx (see Table 2.2). We confirmed the service and the version with the respondents. We agreed with respondents not to share the versions of the software so as not to facilitate attackers. For all the vulnerable services that we observed with the passive scan data, the actual running service and version were the same, according to the respondents. We also confirmed 27 Microsoft services: 1 MS Internet Information Services 8.0 service, 5 MS Internet Information Services 8.5 services, 9 MS Internet Information Services 10 services, and 12 MS HTTPAPI 2.0 services. Again, all versions were confirmed by the respondents. In total, all 42 observations were correct. While the measurements of the versions of the Microsoft systems are correct, the measurements do not provide insights into their actual vulnerability. First, we cannot determine CPEs that can be linked to specific CVEs. Second, we cannot actively verify vulnerabilities by exploiting them, as described in Section 2.4.1. During the interviews, we verified the versions but we did not discuss specific security updates. Table 2.2 provides an overview.

**Table 2.2: Validation of CPEs identified by Censys**

Service	# Correct	# Incorrect
Apache	6	0
OpenSSH	4	0
PowerDNS	3	0
Nginx	2	0
IIS 8	1	0
IIS 8.5	5	0
IIS 10.0	9	0
HTTPAPI 2.0	12	0
<b>Total</b>	<b>42</b>	<b>0</b>

## 2.6. ATTRIBUTION AND UNCLEAR RESPONSIBILITIES

As discussed in Section 2.4.2, our research approach led us to discover a new explanation for the presence of the vulnerable systems: the misalignment between the IP ranges that the municipalities registered with the CERT and the IP ranges that the respondents thought they were responsible for. Perhaps some vulnerable systems persist because the municipalities do not see them as their responsibility, even though they might be notified about their vulnerable status.

In this section, we further explore this misalignment. First, we briefly describe the role of the CERT and the crucial role that the registered IP ranges have in the institutional incident response system. Next, we analyze the relationship between the IPs registered with the CERT and the IPs that the practitioners provided. Then, we quantify the misalignment of the IPs between the CERT and the municipalities. Finally, we examine perspectives on responsibilities.

**Table 2.3: The number of hosts, number of versioned hosts and number of vulnerable hosts that the CERT attributes to a municipality, and the practitioners at the municipality themselves.**

Muni Id	Hosts Muni	Hosts CERT	Versioned Hosts Muni	Versioned Hosts CERT	Vulnerable Hosts Muni Only	Vulnerable Hosts CERT Only	Vulnerable Hosts Shared	Total Vulnerable Hosts	Total Unique CVEs
1	26	0	5	0	3	0	0	3	1
2	3	6	0	3	0	2	0	2	45
3	19	13	11	11	0	0	2	2	3
4	22	30	1	1	0	1	0	1	53
5	11	25	3	8	0	0	4	4	14
6	31	47	2	23	0	4	1	5	88
7	21	12	3	5	1	4	0	5	84
8	18	10	3	6	2	5	0	7	73
9	77	156	26	34	0	0	2	2	37
10	24	16	0	0	0	0	0	0	0
11	24	81	0	0	0	0	0	0	0
12	26	50	1	4	0	1	0	1	57
13	13	7	4	0	0	0	0	0	0
14	113	29	8	6	0	1	0	1	31
15	26	0	2	0	0	0	0	0	0
16	46	1	5	0	0	0	0	0	0
<b>Total</b>	500	483	74	101	6	18	9	33	486

### 2.6.1. ROLE OF CERT AND MUNICIPAL IPS

The IBD-CERT supports municipalities with security advice and liaises between municipalities and the national CERT. On behalf of municipalities, it contributes to the Baseline Informatiebeveiliging Overheid (BIO) – the Dutch compliance framework for information security within government. One of the goals of the CERT is the detection of incidents and crisis situations and the sharing of knowledge between municipalities and suppliers. While municipalities are ultimately responsible for monitoring their own systems, in pursuing its goals, the CERT also monitors the Internet-facing infrastructure of municipalities. In doing so, it also receives information about vulnerable systems from other parties, such as the national CERT, their own scans, ethical hackers, Shodan, and the Dutch Institute for Vulnerability Disclosures (DIVD).

### 2.6.2. MISALIGNED THREAT LANDSCAPE

The misalignment in monitored IP addresses translates to misalignment in the perceived threat landscape by the CERT and the practitioners at the municipality. We find that the CERT generally observes more versioned and vulnerable hosts than the practitioners at the municipality itself. Table 2.3 describes the number of hosts, number of versioned hosts, and number of vulnerable hosts that the CERT attributes to a municipality, and the practitioners at the municipality themselves. It shows that respondents observed 6 vulnerable hosts that the CERT did not observe. The CERT observes 18 vulnerable hosts that the respondents do not. The respondents and CERT both observe 9 vulnerable hosts. In total, there are 33 vulnerable hosts for the municipalities.

The data in Table 2.3 shows that neither organization observes the full set of vulnerable hosts. The CERT observes more vulnerable hosts than the practitioners at the municipalities themselves. This would lead the CERT to send notifications to the respective

municipalities, who, in turn, would not recognize the system. But the CERT also has a blind spot, the vulnerable systems in the public IP range of the municipality that the CERT does not monitor. In those cases, the CERT could not exercise its supportive function, and would not send any notifications at all.

The differences in vulnerable hosts between the organizations impact the CVEs associated with a municipality. We find that the 15 vulnerable hosts at municipalities result in 62 unique CVEs within IP addresses for which they consider themselves responsible. The vulnerable systems observed by the CERT for those municipalities result in a much larger set: 481 unique CVEs. It is not just the sheer number of vulnerable hosts that is larger. We also see a remarkable difference when we look at the CVSS rating of the vulnerabilities (a score between 1 and 6.9 is considered “low/medium”, between 7.0 and 8.9 is “high” and between 9.0 and 10.0 is “critical”) [175]. We find that there are 15 critical CVEs in the IP space identified by the respondents versus 107 critical CVEs in the IP space that are registered with the CERT. Similarly, we find that there are 27 high CVEs in the municipality IP space and 191 highs in the CERT IP space. Table 2.4 depicts the CVEs by severity per organization. The key consequence of the misalignment of the IP ranges is this: the CERT observes much greater risks for the municipalities, compared to the municipalities themselves.

**Table 2.4: Number of vulnerabilities by CVSS3 severity, as observed by each organization. In parentheses are the number of unique IPs with a vulnerability. Note that most IPs run a service with multiple CVEs of different severity levels, skewing the individual and total IP count.**

Org.	Critical	High	Medium/Low	Total
Muni	15 (7)	27 (14)	20 (9)	62 (15)
CERT	107 (18)	191 (27)	183 (22)	481 (27)

This finding can explain two things. First, it means that the CERT – and in its wake, the national CERT and central government agencies – see the municipalities as much more vulnerable than the municipalities’ security practitioners see themselves. That can translate into the perception of the CERT and other government entities that municipalities show a lack of urgency about these issues. Second, it means that the vulnerable hosts persist because the vulnerability notification system is broken. The municipality receives the notifications, but the bulk of these are deemed to fall outside their scope of responsibility. Conversely, there are vulnerable hosts in the self-reported IP ranges that are not registered with the CERT. The municipalities will not be notified about those. Both scenarios lead to vulnerable hosts persisting over time.

During the interviews, most respondents stated that they do not have a complete or up-to-date overview of all the systems that the organization is running externally, so outside the ranges they feel responsible for, but inside the ranges registered with the CERT. Some respondents stated that the internal processes at the municipality for departments to report external systems that they use to the IT or security team were unclear. One respondent (#12) stated “we have a view on those systems that are reported. But not on the ones that are not reported”.

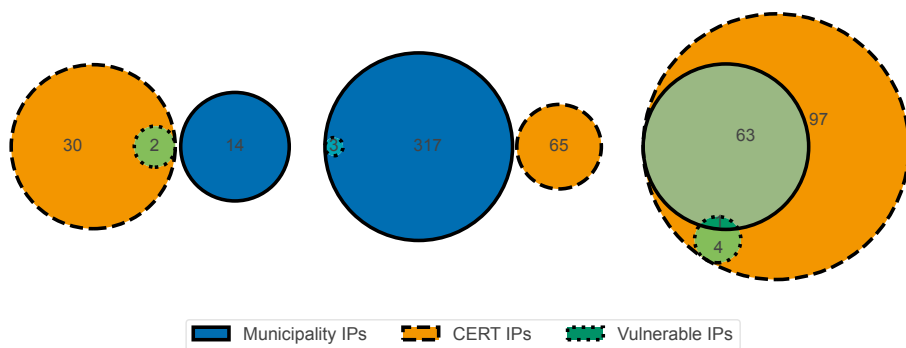


Figure 2.2: Three examples of IP sets registered with the CERT and the IPs used by a municipality.

### 2.6.3. QUANTIFYING THE MISALIGNMENT

To quantify this misalignment, we compare the *sets* of IPs that the CERT associates with a municipality and the IP addresses that the respondents reported as being responsible for. To measure the similarity of those two sets, we use the Jaccard similarity coefficient and the Szymkiewicz-Simpson coefficient, also known as the Overlap coefficient. Both scores range from 0 to 1. The Jaccard score tells us how similar two sets are, where 1 means that all the items appear in both sets. However, the score does not capture when one set is much bigger than the other. Therefore, we also include the Overlap coefficient, which is 1 if one set is fully subsumed in the other set. We find that, on average, the Jaccard coefficient is 0.21 and the Overlap coefficient is 0.56. This means that most municipality IP ranges and CERT IP ranges for that municipality are of a very different size, and also have only partial overlap – meaning, they monitor different addresses. More importantly, the incident response and vulnerability notification infrastructure assume they are 1.

We can also visualize the misalignment using a Venn diagram. Figure 2.2, shows three examples of sets of IP ranges with vulnerable IPs from the CERT and the municipality itself. The remaining diagrams are in Appendix A.4. Clearly, the IP ranges registered with the CERT, and the IP ranges used by our respondents are very different.

### 2.6.4. PERSPECTIVES ON RESPONSIBILITIES

We tried to understand what explains the differences in the attribution of vulnerable hosts. Respondents made a distinction between systems they administer and systems that are used by the organization but are not administered by the respondents, such as SaaS services.

**Administering Systems Themselves.** Most respondents stated that their responsibility was primarily the systems they administer themselves, and exposing them through their public IP range – i.e., the IPs used for routing Internet traffic in and out of their internal network. They minimize their Internet-facing footprint because it provides the administrators with two clear advantages. First, the “front door” is small and easier to man-

age for “administrator”-type duties – it gives them a better overview. As one respondent (#2) stated, “the more external points you have, the harder it becomes and you quickly lose oversight”. Second, the security tools only need to monitor a small set of addresses which reduces the capacity needed for monitoring those systems. One respondent (#5) stated, “Monitoring simply takes a lot of time, and then we are not even taking into account any type of response”. Similar to [7], we found that many municipalities faced staffing issues and limited resources. Consequently, a lot of decision-making prioritized optimized use of (human) resources.

**Not Administering Systems.** The challenge of overseeing systems is exacerbated when external parties provide a service for the municipality. Most respondents observed a trend in the growing number of Software-as-a-Service (SaaS) platforms. They applaud this trend as it reduces the number of applications and servers they need to manage on-premise. This allows practitioners to handle more work with fewer people. However, the downside is that practitioners are not always aware of those systems, and they no longer administer or control the systems on which their data now resides.

Administering systems is difficult for respondents because of the sheer number of business services a municipality offers, each with its own software application. It is this wide array of services that distinguishes it from a regular enterprise according to one respondent: “A municipality has so many types of connections, partners, and disciplines that it’s not comparable to a business. There, they want a single solution but that doesn’t exist in municipalities. It’s very complex. For example, you could get a notification that the tile of a curb is positioned wrong. The team responsible for that has its own application. We have many demands and wishes and we have a great number of running tenders. Unfortunately.”

A majority of respondents stated that there was no clear overview of the platforms that the municipality uses, as those platforms were not always reported to the IT team. For some respondents, a clear process to register newly connected platforms or a mandate to enforce the registration obligation at the IT team was lacking. This is the well-known problem of “shadow IT”, which has been plaguing IT managers for decades now [204]. In only two cases did the respondents know the vulnerable host at the IP address that we shared with them, even though they were not responsible for administering it and it wasn’t in their public IP range. They knew it was running at a third party. Respondents also brought up the issue of control. They indicated that while procedures are in place to demand security measures during the tender, they do not have the power to exert control. For example, a respondent stated that monitoring – i.e., scanning – external systems was frequently not allowed by the external vendor. Another respondent mentioned that the municipality wanted to ingest logs from the vendor into their SIEM, but the vendor did not want to share its logs, one of the voiced concerns was the privacy of (non-municipality) user information. In all of those cases, the practitioners could not exercise control over how their data was managed at the partner organization, yet they feared that the (political) fallout in case of an incident would be their responsibility.

**Administering Shared Resources.** Some municipalities collaborate in a governance structure to share ICT resources. This governance structure streamlines resources to re-

**Table 2.5: Features of the municipalities, the number of total vulnerable hosts, and the Jaccard and Overlap indexes with CERT. Most municipalities collaborate and provide IT services for several municipalities and auxiliary organizations.**

Muni ID	Vulnerable Hosts	Inhabitants	IT team	Security Team	Servers	Muni's	Aux orgs	Jaccard	Overlap
1	3	10-100k	10-20	yes	100-200	1-2	0	0	0
2	2	10-100k	1-10	yes	0-100	1-2	0	0	0
3	2	100-200k	20-30	no	100-200	5-10	1-2	0.36	0.56
4	1	10-100k	1-10	yes	100-200	2-5	>3	0.038	1
5	4	10-100k	1-10	no	100-200	1-2	0	0.77	1
6	5	200-300k	30-40	yes	>1,000	>10	1-2	0.38	1
7	6	100-200k	20-30	yes	500-1,000	1-2	0	0.16	0.46
8	7	100-200k	10-20	yes	100-200	1-2	0	0	0
9	2	300-400k	30-40	yes	>1,000	5-10	1-2	0.98	0.99
10	0	200-300k	30-40	yes	200-500	1-2	0	0.17	0.79
11	0	100-200k	30-40	no	500-1,000	1-2	0	0.27	1
12	1	10-100k	1-10	yes	100-200	3-5	0	0.47	1
13	0	200-300k	40-50	no	>1,000	3-5	0	0.10	0.18
14	1	100-200k	20-30	yes	200-500	1-2	0	0.02	0.96
15	0	200-300k	40-50	yes	>1,000	5-10	1-2	0	0
16	0	100-200k	40-50	yes	>1,000	1-2	1-2	0	0

duce costs and optimize the capacity of scarce IT personnel. However, it is not without problems. Respondents in such a governance structure stated that it is hard to draw the line between responsibilities. For example, in two interviews, we spoke to respondents about where the line was drawn between the hardware layer and the application layer. The idea behind this distinction is that the municipalities themselves can manage much of the application layer to accommodate their specific business needs. However, in practice, these layers are often intertwined, and it becomes unclear who is responsible for mitigating a vulnerable system. We observed one case where the organization that runs the infrastructure knows about vulnerable systems in the application layer but is unable to patch the system, mainly because they cannot oversee the impact of an update on the systems that provide the business services.

## 2.7. MANAGING VULNERABLE SYSTEMS

The last explanation for exposed municipal hosts looks at the reasons that organizations might not patch a system they are responsible for. Practitioners may be unaware of vulnerable systems, or they might be unable to patch them, or they might have decided that patching is not needed, e.g., because they have specific mitigation strategies in place, such as firewall rules or monitoring.

We identified 15 vulnerable systems for which the practitioners we interviewed considered themselves responsible. We asked these respondents how they dealt with those systems. We asked for their rationale on patching, why the systems were not patched, and what if any, other actions had been taken. Due to the nature of our collaboration, we did not learn of the specifics of the business service that a vulnerable system provided. That said, in general, practitioners mentioned that most of their systems contain valuable data, such as personally identifiable information, albeit frequently fragmented. Every bit of valuable data or provided service is considered important, and compromise would have privacy implications even if it applied to only a handful of citizens.

### 2.7.1. IDENTIFYING VULNERABLE SYSTEMS

We asked respondents how they identified vulnerable systems in their daily work. All respondents stated that they use vulnerability scanners to do so. In addition, they stated that yearly penetration tests are conducted. To stay up to date on the latest vulnerabilities (that might not be incorporated in the vulnerability scanner), they receive security advisories from the CERT, vendors, and popular security news sources. Our findings on the ingress of vulnerability information are in line with earlier work [54], in that practitioners relied on curated vulnerability information from authoritative sources to consider vulnerability information. As we will discuss below, for some vulnerable systems, sysadmins did not know the version of the software they were running. So, they relied on external triggers to become aware of vulnerabilities.

There is a compliance framework in place to act on security advisories. As stated in Section 2.6.1, the Dutch government uses a compliance framework, BIO, to improve and measure security practices. It contains a chapter on vulnerability management which states that if the severity level of security advisory of the national CERT is marked with probability as “High” and impact as “High”, (also known as a “High/High”), the vulnerability should be resolved or mitigated as soon as possible and at the latest within a week. All respondents noted that if they received a “High/High” advisory, they would move to action almost instantaneously.

### 2.7.2. VULNERABLE VS. NO VULNERABLE SYSTEMS

We analyzed the interviews to determine if municipalities that did not have vulnerable hosts did anything differently than those with vulnerable hosts. We included 5 municipal IT organizations where we detected no vulnerable hosts. Remember that our measurement approach relies on obtaining versions from banner data. If administrators hide the version information, our method will not determine vulnerable hosts. We learned that 2 of the 5 municipal organizations indeed ran security products that obscured versions.

At 11 municipal organizations, we did observe vulnerable hosts, either in the IPs registered at the CERT or administered by the municipality. We tried to find a common denominator for organizations with vulnerable hosts versus those without. We looked at security tools, the capacity of IT staff, the size of the municipality, the size of the IT team, the presence of a security team, and the act of vulnerability scanning.

All the respondents we interviewed, with and without vulnerable hosts, had basic generic security tools in place, like firewalls, NAT, EDR, logging, and network segmentation. Similarly, all respondents noted that they did not have sufficient capacity, in terms of qualified IT staff, for their security duties. Next, we checked if the size of the municipality, measured by the number of inhabitants, had a relation to the number of vulnerable hosts. We observed that both the smallest and largest municipalities had vulnerable hosts, as did several in between. We then compared the size of the IT team but also found that vulnerable hosts occurred in small and bigger IT teams. Finally, we looked at organizations in terms of the number of servers they ran. These servers are not all exposed to the Internet, but the number of app servers acts as a proxy for the size of their infrastructure. The idea is that a larger infrastructure might contain more vulnerable machines. However, here too we see no clear distinction. Next, we observed that all respondents, except one, engaged in vulnerability scanning, discounting this as an explanation for the

presence or absence of vulnerable hosts. The one respondent who did not actively use it was not part of the security team. Finally, we wondered if the presence of a security team might have an effect. From the 16 interviews, 12 organizations ran a security team, and 4 did not. We observed that 9 municipalities with a security team have vulnerable hosts. Of the 4 municipalities without a security team, 2 municipalities exposed vulnerable hosts.

In sum, we do not observe a link between specific features of an organization and the number of vulnerable hosts. We also do not find substantially different security practices among respondents in the interview data.

### 2.7.3. NO PATCH

We hypothesized that observed vulnerable systems are indeed vulnerable, but administrators may have their reasons for not patching. We first examined if they had put specific mitigation strategies in place for the vulnerable hosts. All the respondents stated that they had generic mitigation strategies in place (such as network segmentation, logging, firewalls, and Intrusion Detection Systems). Some respondents stated that they also use a managed SIEM. However, only in one case was a specific mitigation strategy in place for an observed vulnerable system. That system was run on an isolated network. Without specific mitigation strategies in place, we asked administrators what other actions, if any, were taken for the vulnerable hosts. We analyzed the interviews for the 15 vulnerable hosts and synthesized the responses into various explanations that we condensed into four categories: *unaware* of the vulnerable system, *unable* to patch it, the system was (in the process of being) *retired*, or the system was *shut down*. We describe each in more detail below, and tabulate an overview of these explanations in Table 2.6.

Table 2.6: Explanations for vulnerable systems

Explanation	# Systems
Unaware	9
Unable	3
Retired	3
Shut down	N/A
<b>Total</b>	<b>15</b>

**Unaware.** We observed during the interviews that respondents were not always aware of the vulnerable system. We encountered three types of unawareness. First, an administrator retired the system and assumed it was no longer online. This, however, was not the case. One respondent (#6) said: “this is a system that is phased out. I’m actually surprised about this. Thanks, I’ve got some homework to do.” Second, there was a case where a vulnerable system was not on the radar at the organization. When presented with the system, the security team could not find the system in their asset inventory but acknowledged it was running in their IP range. There was no direct explanation for that situation. Third, the system was not directly identified as a vulnerable system. The respondent (#1) stated: “These servers run directly from the Debian repos. We did not patch these systems because we assume the repo provides a decent package”. For this

particular case, the administrator was not part of the security team, so it may have been flagged as vulnerable elsewhere in the organization. However, the actual administrator of the system initially did not consider it vulnerable.

We investigated this type of unawareness further by checking if respondents knew about the versions of the software they run and if they were vulnerable. None of the respondents during the interview explicitly had the versions of the software they were running at the top of their minds. Similarly, none of the 29 respondents directly knew the latest version of the software they were running. All had to look up the service and version from their asset inventory system, most often a vulnerability scanning report.

We tried to gauge if respondents knew *ex-ante* if they set up a vulnerable system in their infrastructure. Most respondents stated that the software that is run is installed to the latest version when set up, regardless of potential vulnerabilities, because that is the best they can do. Potential vulnerabilities will be found when the vulnerability scanner is run. One respondent indicated that when installing the latest version, he checked the latest packages for information and vulnerabilities from the distributor's repository to make sure a newer version would not be released in the very near future.

**Unable.** In two cases, practitioners reported that they were unable to patch a vulnerable system. This happened either because of a lack of mandate from the organization or because the vulnerable software was a dependency in a product that was used to provide a business service.

At one organization, the security team was aware of the vulnerable system but was unable to undertake mitigating actions. This particular situation derived from a governance structure where the respondents were part of an organization that was responsible for the majority ICT infrastructure of the municipalities but not the last application layer. The application layer was managed by a small IT team at the municipality itself. The security team at the organization managing the infrastructure is somewhat involved in the management of the application layer due to their expertise, but they did not have the mandate to intervene themselves. Another consideration was that a patch could break the services offered by the vulnerable system. A respondent (#13) stated “fixing vulnerabilities in the application layer is outside the scope of our mandate. That said, it isn't entirely that black and white. But in this case, we cannot functionally oversee the consequences for the underlying application and business processes.”

At another organization, the vulnerable host ran a product with a dependency that was vulnerable. The vulnerable software could only be patched by updates from the product. The risks of that system were mitigated by the monitoring of an endpoint security product – that ran on all managed devices – and by running the system on an isolated segmented network. For that host, the security team did not directly intervene in mitigating the risk but had to engage other teams to act on the vulnerable system. The respondent (#12) stated: “that's why I'm pushing these people to act on this system”. But sometimes such a system could not be removed. The respondent (#12) said: “this particular system, it is provided by a supplier and someone in our organization opted for that product. How do you deal with it? Retiring the system is the only way”. The system was not actually retired because of business reasons.

**Retired.** In various cases, the observed vulnerable hosts were in the process of being retired. A retired system should not be online, but there is some time between deciding to retire the machine and it actually being offline.

Retiring a system could have various reasons but, most often, respondents stated that it was a (legacy) system that is outdated. For example, some systems ran outdated software, and the service it provided can now be done better with a new product. Consequently, the system was phased out and didn't get much attention anymore. The legacy system was marked for retirement and lingered around for a while before it actually went offline. A respondent (#6) said "this was a great product at the time but these days it's outdated. We are now in the phase of retiring this system". In another case, we found a vulnerable system before the interview, and during the interview, the system was no longer online, the respondent (#23) stated that: "by now that system is turned off". This happened only once.

**Shut down.** One of the most drastic mitigation strategies for vulnerable systems is simply "pulling the plug". This was not done for any of the vulnerable systems we observed, but many respondents stated that shutting a system down was part of their toolbox of mitigation strategies – albeit one very few actually want to use. When asked if this was actually done in practice, only a few respondents stated that this was within their power to actually execute. The majority of respondents stated that, while it is an option, in practice, shutting down is hardly done because business continuity takes priority. For most respondents, the only situation where they actually shut down systems was during the Log4J vulnerability. The severity of the vulnerability and the fact that it was not clear which software was vulnerable allowed for enough organizational pressure to trump business continuity. As one respondent stated: "once a vulnerability hits the news, people start taking it seriously. Sometimes we need an incident like that to make strides in security".

But without that sense of urgency, business overpowers security. One respondent (#22) stated, "business and security sometimes have opposing interests. Contrary to a commercial organization, a municipality has certain societal obligations, therefore we simply can't shut a system down like that because we are required to offer those services. Considerations are complex and discussions are quickly taken out of context. Then it doesn't matter what actually happened, but how it is perceived because something is in the newspaper".

#### 2.7.4. PATCHING SYSTEMS: PRIORITIZING

Respondents stated that patching is the preferred strategy to deal with a vulnerable system. Sometimes, a patch is not directly available, as was the case with a Citrix vulnerability in 2020 [2], then they rely on the mitigation strategies provided by the vendor or the security advisory. However, the capacity, in terms of people, for rolling out patches is limited. The vast majority of respondents stated capacity as an obstacle to security. As practitioners face vulnerabilities in their Internet-facing systems as well as their internal networks, this forces them to prioritize what systems to patch first. In doing so, there are two main criteria: a) whether the system is Internet-facing and b) the severity level of the vulnerability.

First, several respondents stated that all vulnerabilities should be dealt with. This self-reporting, however, is contradicted by the fact that we did find vulnerable hosts, illustrating the limitations of self-reporting on patching behavior. To be fair, many respondents also appeared to accept that there will always be vulnerabilities somewhere in their infrastructure. Vulnerabilities in Internet-facing systems should be dealt with first – as they consider it the front door to their internal network. One respondent (#26) stated: “Internet-facing systems have priority because the chance of abuse is higher than systems inside our network.” If this is true, then there should be more vulnerabilities in internal systems versus the Internet-facing systems. We could not measure this directly, but when asked, several respondents admitted that there were indeed (many) more vulnerabilities in their internal network.

While many respondents stated that their Internet-facing systems should not contain vulnerabilities, it was also mentioned that they do not consider it likely that a real attack would happen there. Instead, they feared an attack via an unsuspecting user clicking a link in a phishing mail.

Second, respondents indicated that systems with a high-severity vulnerability are prioritized. The severity metric is most often determined by either the severity level of the security advisory or the severity score of the vulnerability scanner. The Common Vulnerability Scoring System (CVSS) cite[75] is a popular scoring system to determine the severity of a vulnerability. The CVSS score is popular but has its limitations. For example, it does not take into account the ease of exploitability of a vulnerability, the availability of an exploit, or details on the number of exploitations of the vulnerability in the wild. Many security companies expand on CVSS with their own data and ranking to improve the assessment of the severity of a vulnerability. In doing so, (proprietary) vulnerability scanners often report the CVSS as well as their own scoring system. For example, the vulnerability scanner Nessus – popular among respondents – provides its own Vulnerability Priority Rating (VPR). The VPR takes additional factors into account, such as the CVSSV3 Impact score, the age of the vulnerability, the exploit code maturity, and more [239]. If vulnerabilities have a High or above classification (in any kind of scoring methodology), they are quickly prioritized according to the respondents.

To verify this claim for high-severity security advisories, we examined the national CERT security advisories for “High Impact/High Probability” vulnerabilities and referenced them with the systems we scanned. The security advisories contain vulnerabilities for open-source software but also proprietary software. We could not validate the claim that respondents patched “High/High” advisories for proprietary software quickly because the associated software services could not be fingerprinted by us for a version. However, in one case, security company Fox-IT wrote in their blog that they could fingerprint Citrix software for two specific CVEs [84]. In the CERT dataset with IP addresses, we reproduced their method and observed that the systems with CVEs were applicable for 7 hosts. Looking historically, in the two weeks after the advisory was sent and the update became available, these systems were patched, giving credibility to the respondents’ claims.

## 2.8. DISCUSSION

**Explanations for persisting vulnerable systems.** We consider three explanations for the 154 vulnerable systems we observed in the IP ranges registered with the CERT: incorrect measurements, misalignment in the responsibility of IPs, and vulnerable machines unpatched for some other reason. We interviewed 16 municipal IT organizations that had 33 of the vulnerable systems.

We learn that the first explanation doesn't explain any of the vulnerable systems. We find that the external measurements of hosts using banner information are not plagued by false positives. The second explanation, the misalignment of IPs – i.e., IPs registered at the CERT and those used by practitioners – is observed at all the 16 municipal IT organizations we spoke, and this most likely happens at many more, if not all, municipalities. Of the 33 vulnerable machines for the municipalities, 18 vulnerable hosts were seen only by the CERT, 6 by the municipalities alone, and 9 were seen by both organizations. This brings us to the third explanation: vulnerable machines that are unpatched for other reasons. Of the 15 vulnerable hosts observed by the municipalities, 9 are explained by administrators being unaware of those systems.

In short, the main explanation is the misalignment in the attribution of IP ranges, where administrators do not consider the systems their responsibility. This explanation is followed at some distance by the explanation that the IT organization was unaware of the presence of the vulnerable systems.

What is causing the misalignment problem? As Vermeer et al. [262] noted, organizations consistently struggle to keep a complete inventory of their assets. The assets are constantly changing, with many changes unplanned or unrecorded. This is closely related to the problem of “shadow IT” – systems and services that are “not known, accepted and supported” by an organization's official IT department [204]. Indeed, when respondents were speculating what the vulnerable hosts were, they frequently mentioned SaaS solutions and specific services contracted by some department of the municipality, but outside their purview, the purview of the IT department. This is classic “shadow IT”. This explains why they did not consider it their responsibility to safeguard these systems. It also suggests that this is most likely not a problem exclusively to municipalities, and we might expect this also to occur in other enterprise environments. Not only does it mean there is no clear responsibility to keep those systems secure, but it also gravely undermines the CERT-based notification mechanism. Those notifications reach the IT department, but cannot find their way to the actual entity managing the host, because IT does not know. The fact that the IT teams see themselves as powerless towards “shadow IT” does not reduce the actual risk for the organizations. The vulnerable systems continue to run, exposing 183 “Low/Medium”, 191 “High”, and 107 “Critical” CVEs, according to IPs from the CERT data.

Finally, our findings also highlight a discrepancy between the widely-held view that local governments are very vulnerable, as mentioned in the Introduction, and the perspective of the local IT departments, as the latter observe far fewer problems in their own systems. This discrepancy might explain why the current situation persists, even though CERTs and higher levels of government keep warning local governments.

The municipal IT departments do their work under serious resource constraints. The lack of capacity and staff frequently came up. We observed that vulnerable hosts, with

the exception of one, did not have specific mitigation strategies in place. Instead, respondents rely on generic mitigation measures. Respondents frequently mention the lack of capacity for IT security tasks. This aligns with our finding that organizations have only generic security measures in place – a rational choice to maximize defenses with limited resources. Coincidentally, in December 2022, the Association of Dutch Municipalities (VNG) requested additional resources from the central government to increase its IT capacity in support of the Dutch National Cybersecurity Strategy, but those resources were denied [130]. Simply put, cybersecurity requirements increase, yet resources to comply lag behind.

**Recommendations.** So, what can municipalities do to tackle the attribution issue? After all, they themselves registered the IP addresses with the CERT. Clearly, they would benefit from keeping the IP ranges registered with the CERT up-to-date. At a minimum, this allows notifications to reach the correct entity. If they are unaware of vulnerable systems, the notifications should help address that. What appears needed is some guidance or support on how to handle the responsibility gap that exists around “shadow IT”. Who is responsible for what system? The owner of the system? The owner of the data in the system? Respondents occasionally stated that they feel somewhat responsible for systems running elsewhere since they handle municipal data. At the same time, another respondent stated that they are not allowed to scan the systems of their partners, so they are unaware that their data resides in a vulnerable system. Until a clear delineation of responsibilities exists, organizations remain at risk for cyber threats. The new NIST 2.0 framework may guide practitioners in this respect. It emphasizes a new Govern function to gain a better “understanding of cybersecurity roles and responsibilities” [173].

For CERTs, our findings suggest two points. First, the institutional framework to mitigate cyber threats assumes that vulnerable systems are correctly attributed. This assumption turns out to be problematic. The national CERT detects a vulnerable system and notifies the relevant sectoral CERT. In turn, the sectoral CERT notifies the appropriate constituent about the vulnerable system. The constituent is then expected to take action. But in practice, we observe that the last, crucial step of this notification process is flawed – as the owner of the vulnerable system is often not the recipient of the notification. As a result, we find that those systems remain vulnerable, with the data of the organization at risk. This finding contests the effectiveness of the institutional framework for notifying vulnerable entities and highlights the need for a better reporting process between the CERT and its constituents.

Second, for CERTs, like for municipalities, there seems to be a need for a clear delineation of responsibilities for external systems. Currently, what constitutes the network of an organization appears to be diffuse, but counting only on-premise infrastructure seems archaic. Perhaps a recommendation could be to build a feedback loop for the notifications, a bit like a ticketing system. This way, if a notification is not picked up by the entity to whom it is assigned, because that entity sees it as outside its responsibility, then the ticket gets returned to the CERT. Both the CERT and the municipal IT leadership, e.g., the CISO, can then observe what systems are vulnerable, yet not acted upon. This is then a starting point for identifying who is managing those systems. Currently, neither the CERT nor the municipalities seems aware of the scale problem we have uncovered.

If notifications function like tickets, then the scale and location of the problem become very clear to see for all parties involved.

2

**Limitations.** Our research design introduces several limitations: external validity, internal validity, scope, and desirability bias. First, we focused exclusively on Dutch municipalities. This risks that our findings might not be generalizable. We believe, though, that CERTs worldwide face similar challenges in delineating responsibilities and correctly attributing vulnerable hosts to their constituents for effective notifications and subsequent remediation. Our findings point to the problem of “shadow IT”: the challenges of managing asset inventory, including IPs and externally run services, which is a problem that many other organizations deal with. Also, a sample size of 16 interviews is limited. This sample, however, does cover 17% of the total population of municipalities. Second, our respondents were directly involved with managing vulnerable systems. Yet, they are only part of the (security) IT teams within the organizations, and their knowledge might be incomplete. We allowed additional respondents during the interviews to mitigate this risk. Yet, the limited knowledge is not just a barrier to measurement, but also an operational reality with consequences: respondents did not know who was responsible for large portions of the IP ranges that were registered with the CERT. Hence, they cannot delegate vulnerability notifications, let alone ensure mitigation.

Third, we could only discuss Internet-facing systems that ran a service with a version. Most proprietary (security) products run a service that does not include (backported) version information in the banner. Thus, the number of exposed vulnerable hosts is likely to be higher. That said, our methodology is limited as we could not perform active measurements of observed vulnerabilities. Therefore, the measured vulnerability of systems does not directly map to an equal amount of risk.

Lastly, the interviews risk desirability bias. When asked why systems were vulnerable, participants might give answers to paint them in a favorable light. We tried to mitigate this risk by a) talking about actual systems, b) interviewing respondents without their superiors, c) keeping the interviews confidential for the municipality and the CERT, and d) stating in the interview that we were scientists and were not there to pass judgment. We also believe that our methodology, interviewing respondents about actual systems, led us to a new discovery of the misalignment in responsibilities for vulnerable systems. Without our approach, practitioners would have reported very few, if any, vulnerable systems in their networks.

## 2.9. CONCLUSION

We asked what explains unpatched vulnerable systems that are detected and notified about, but not patched. We found that the detections are correct. It turned out that most of the vulnerable systems fell into a responsibility gap around “shadow IT” and other systems outside the reach of the IT organizations. Sysadmins did not consider these systems their responsibility. The CERT was not aware. We further identified that for most systems that did fall under the sysadmin’s responsibility, they were unaware of the presence of the vulnerabilities. Our findings highlight the need to re-evaluate and improve the critical institutional structures of incident response and vulnerability notifications.

# 3

## “TELL THEM THEY ARE A RESPONSIBLE ENTITY, NOT A CUSTOMER”: UNDERSTANDING PRACTITIONER CHALLENGES IN SECTOR CSIRTS

*In this paper, we study the experiences of practitioners in sectoral Computer Security Incident Response Teams (CSIRTS)—specialized teams that mediate between national cybersecurity authorities and the sector constituency. Through interviews with 18 professionals connected to the Informatiebeveiligingsdienst (IBD-CSIRT) for Dutch local governments, we uncover tensions in how key services are valued. For vulnerability notifications, while the CSIRT staff consider them a core service, many constituents hardly mention them, and systemic gaps in information forwarding mean that crucial alerts often never arrive. We extend these insights with 5 interviews across other sector CSIRTS and a validation workshop with 7 participants, all security officers from sector CSIRTS, revealing shared challenges in balancing technical expertise with sector knowledge, building trust-based relationships, and navigating institutional bottlenecks. Our findings contribute the first systematic account of how sector CSIRT professionals understand and perform their role, highlighting the tensions in providing sector-wide support to professionals with differing security needs.*

---

This chapter has been published as: **Ethembaoglu, A.M.**, Kadenko, N.I., & Angelova, Y., & Zhauniarovich, Y. & Parkin, S. & van Eeten, M.J.G. (2026). ““Tell Them They Are a Responsible Entity, Not a Customer”: Understanding Practitioner Challenges in Sector CSIRTS”. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*

### 3.1. INTRODUCTION

Computer Security Incident Response Teams (CSIRTS) are at the front line in responding to cybersecurity incidents and attacks. They consist of specialized security professionals responsible for handling, responding to, and preventing cybersecurity incidents such as data breaches or malware attacks. Enterprise CSIRTS operate within a specific company to protect its internal systems and data.

In addition, most countries also have a national CSIRT. They deliver a wide range of services, such as issuing advisories, sharing threat intelligence, exchanging information with international partners like other national CSIRTS, and notify responsible parties when scans have detected that their internet-facing systems have vulnerabilities. It is critical that these notifications reach the entities responsible for these systems so that they can remediate the problem.

A national CSIRT is a national coordinating body far removed from individual organizations. Over the last decade, this institutional distance has encouraged the development of more specialised CSIRTS closer to organizations in specific sectors: sector CSIRTS. On paper, they perform similar functions to the national CSIRT, but for a more targeted set of constituents within a specific sector (e.g., finance, energy, or water). Institutionally, sector CSIRTS are a model that has been copied worldwide[180]. Practitioners in sector CSIRTS are expected to maintain relationships with their constituents so they can channel vulnerability notifications and other information from the national CSIRT to the relevant affected entities, and to support them with incident response, advisories, intelligence sharing, and sector-specific expertise. Doing so relies not only on network-scanning technologies, but also on the working relationships and sector knowledge that connect national bodies, sector CSIRTS, and heterogeneous local organizations.

Sector CSIRTS have become critical institutions in their own right. In 2016, the US Presidential Policy Directive 41 (PPD-41) highlighted the importance of sector-specific cybersecurity measures and the collaboration between the Department of Homeland Security (DHS) and sector-specific CSIRTS [185].<sup>1</sup> Around the same time, in the European Union, the Directive on Security of Network and Information Systems (NIS) mandated that certain sectors had to establish incident response capabilities, which included sector-specific CSIRTS [67]. Its 2022 successor in the EU, NIS2, strengthens the role of sector CSIRTS by expanding the number of sectors that are required to set up a sector CSIRT from 7 to 18.

Despite this institutional importance, we know little about how sector CSIRTS actually function as socio-technical arrangements in practice: how practitioners in these teams prioritize between incident response, community education, intelligence sharing, vulnerability notifications, and other services; how constituents experience these services; and how governing bodies and national CSIRTS shape what is expected from sector CSIRTS. These questions are not only about technical capabilities, but about how scarce resources are allocated, how legitimacy is built, and how dependencies between organizations affect the services that are ultimately delivered. Currently, the human factor in sector CSIRT teams is poorly understood. Only by understanding the practitioners and practices on the ground can we establish whether and how security may be ad-

<sup>1</sup>The nomenclature can differ per country and in the U.S. this role is fulfilled by ISACs (Information Sharing and Analysis Centers).

vanced [155].

Industry guidance for CSIRTs acknowledges sector CSIRTs but offers little help on these day-to-day challenges. While there is ample guidance for practitioners at ‘regular’ CSIRTs, for sector CSIRT practitioners there is only guidance to get started [180], not on how to operate one. Here, by “operate” we mean the practical work of selecting and shaping services within resource constraints, coordinating with national CSIRTs, and engaging a diverse constituency with different levels of security maturity. The current version of the authoritative FIRST (Forum of Incident Response and Security Teams) CSIRT framework recognizes sector CSIRTs as “special types of CSIRTs” and states that it will describe them in future versions of the framework [77].

In academic work, sector CSIRTs remain uncharted territory. Many studies focus on enterprise CSIRTs, e.g., [73, 99, 124, 142, 162, 210]. The few papers that do focus on sector CSIRTs did not study operational sector CSIRTs, but primarily argue that specific sectors would benefit from setting up a sector CSIRT in light of the threats the sector faces [97, 166, 267]. One exception is a case study of the communication channels used by the Norwegian local government sector CSIRT and its members [227]. In sum, we lack empirical insights into how sector CSIRTs professionals and their surrounding stakeholders navigate the challenges of operating these organizations and providing services in practice.

In this paper, we address this gap with a mixed-method study of practitioners and stakeholders engaged with sector CSIRTs in the Netherlands. Rather than only interviewing sector CSIRT staff, we deliberately study the ecosystem around a sector CSIRT: its practitioners, its constituents, its governing bodies, and the national CSIRT that depends on it to reach the sector. The first phase of our study is a detailed case study of one specific sector CSIRT, the Informatiebeveiligingsdienst (IBD-CSIRT), which serves the local governments sector in the Netherlands. We conducted interviews with 18 professionals who interact with the IBD-CSIRT: staff of the sector CSIRT itself, representatives from municipalities (constituents), and representatives of governing bodies and the national CSIRT. This design allows us to contrast expectations, perceived value, and tensions across stakeholder groups, rather than foregrounding only the “supply side” of sector CSIRTs.

Across these interviews, participants consistently highlighted incident response, community education and expert insights, and intelligence sharing as key services that make sector CSIRTs valuable to their communities. At the same time, a surprising tension emerged around vulnerability notifications. IBD-CSIRT staff understand notifications as a core preventive service and as an important way for sector CSIRTs to add value on top of national CSIRT scanning, yet many municipal participants hardly mentioned notifications at all. This was notable because, on paper, vulnerability notifications are a canonical example of the mediating role of sector CSIRTs: they depend on up-to-date asset inventories from constituents and on systematic forwarding of scan results from the national CSIRT.

To understand this disconnect, we use vulnerability notifications as an in-depth analytic lens on the institutional dependencies and missing feedback loops that shape sector CSIRT work. In the second phase of our study, we therefore complement the interviews with a historical analysis of the vulnerability notifications that were supposed

to be sent to IBD-CSIRT constituents between 2015–2024. Our analysis uncovers a systematic problem in the national notification mechanism: many of the notifications that should have been sent to the IBD-CSIRT, and thus to municipalities, never arrived because the national CSIRT did not forward them. We conducted additional interviews with the professionals involved in the notification program to corroborate this problem and understand its causes. This focus on vulnerability notifications does not imply that they are the only or most important service; instead, they offer a tractable case where we can combine qualitative accounts with log data to reveal how dependencies between organizations can quietly undermine a service that is central in policy documents but nearly invisible to many constituents.

Our in-depth ‘vertical’ approach around a single sector CSIRT is arguably not scalable across many sectors, let alone different countries. So in the third and final phase, we took a ‘horizontal’ approach and interviewed professionals of all-but-one other sector CSIRTS in the country ( $n = 5$ ). We also conducted a validation workshop with most sector CSIRTS in the country ( $n = 7$  workshop participants). This allows us to examine to what extent the service portfolio, tensions, and notification challenges identified in the local-government case generalize to other sectors, and to refine our account of cross-cutting practitioner challenges.

We aim to answer the following two research questions: (i) What are service-specific challenges and expectations of stakeholders (sector CSIRT staff, governance bodies, and constituents) on the services provided by CSIRT practitioners? (ii) What are the strategic challenges for practitioners of a sector CSIRT in providing these services? In sum, we make the following contributions:

- We present the first empirical mixed-methods study on sector CSIRTS. We describe the perspectives of stakeholders and the challenges of practitioners in providing key sector CSIRT services. We find that practitioners’ daily practices are shaped by three dynamics: resources, legitimacy, and dependency. For example, we find that the actual provided services often do not align with the expectations of constituents, especially around incident response.
- We identify several strategic challenges for sector CSIRT practitioners: a diverse constituent population in terms of maturity, trust issues, and service-specific organizational dependencies.
- We evaluate the functioning of the national vulnerability notification mechanism – a key service of the sector CSIRT. We find that many notifications are not arriving at constituents. The practitioners did not detect this problem, signalling a missing feedback mechanism – which is missing almost everywhere in vulnerability notification mechanisms.
- We provide an empirical basis and recommendations for the development of guidelines for practitioners operating in sector CSIRTS.

### 3.2. RELATED WORK

**Frameworks and Industry.** The Forum of Incident Response and Security Teams (FIRST) organization provides leading industry guidance on Computer Incident Response Teams

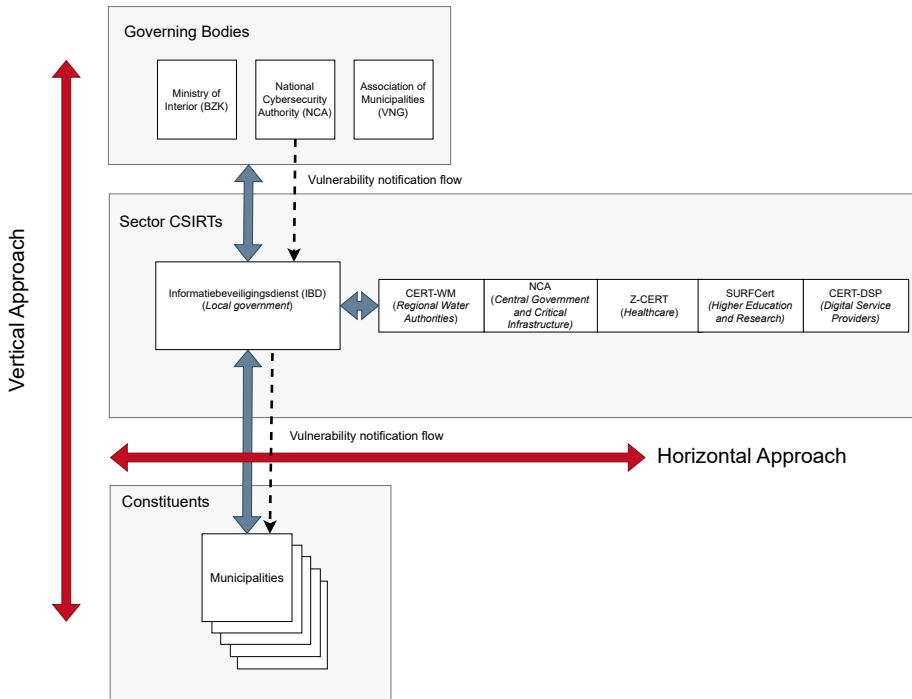
(CSIRTs) [76, 77]. The Computer Emergency Response Team (CERT) Division of Carnegie Mellon University developed a framework to establish sector CSIRTs [180] but does not address the challenges of operating one. Other industry efforts have largely focused on non-sectoral CSIRTs [13, 64–66, 230], guiding the establishment and operation of CSIRTs, as exemplified by the continuously updated handbook for CSIRT teams [112, 271].

**CSIRTs.** Academic research explored various aspects of CSIRTs. Most studies concentrated on enterprise CSIRTs. These studies examined the communication needs, tools, and technical infrastructure required for effective CSIRT functioning [73, 99, 124, 142, 162, 210]. By contrast, limited work has focused on the specific needs and operations of sector CSIRT practitioners. Several studies argue for the benefit of a sector CSIRT in the light of new threats those sectors face [97, 166, 267]. One study described the communication channels used between the Norwegian local government sector CSIRT and its members [227].

**Human Factors in incident response.** Another line of research explored human factors that play a significant role in the effectiveness of incident response teams. Research has demonstrated that the success of these teams depends not only on technical capabilities but also on the dynamics of individuals working together [35, 200, 206, 228, 232, 255]. Similarly, another line of work studies practitioners at Security Operations Centers (SOCs), how practitioners prioritize, conduct assessments, and set up a SOC [1, 38, 219]. In [14], the authors identified challenges and coping strategies for threathunting practitioners. Closely related, there exist additional works that study security workers [8, 52, 141]. These studies noted that, among other things, for many organizations, recruiting and retaining skilled personnel is a challenge.

**Vulnerability Notifications.** Finally, there is a large body of work on (vulnerability) notifications and Coordinated Vulnerability Disclosures (CVD) [101, 192]. In [60, 134, 231] the authors find that notifications lead to higher patch rates. This supports the idea of setting up and running a structural notification mechanism. Similarly, in [252], the authors find that security notifications lead to higher fix rates, compared to privacy notifications. They also identify challenges in reaching responsible parties. Similarly, Cetin et al. showed that retrieving contact information at scale was problematic. But once contacted, entities were more likely to remediate [34]. In [280], the authors found that detailed abuse reports increased cleanup rates, but sender reputation was not important. Li et al. found that by addressing owners of resources directly, vulnerability notifications promoted faster remediation than notifications sent to national CERTs [133]. In [257], though, the authors find that many notifications are not acted on, and the sending entity is unaware of this inaction. These studies all highlight the difficulties in effective notification mechanisms without a feedback loop.

Existing work examined enterprise CSIRTs, SOC practitioners, and vulnerability notification mechanisms, but sector CSIRTs and their role in the broader cyber ecosystem remain largely unstudied. This paper offers the first systematic account of how sector CSIRT professionals understand and enact their role, illuminating tensions inherent in supporting a diverse constituency with uneven security needs. We situate these insights within prior research on practitioner dynamics in SOCs, enterprise CSIRTs, and organizations that consume security services, e.g., the National Vulnerability Database (NVD). In addition, we empirically analyze the sector's vulnerability-notification pipeline and



**Figure 3.1:** Overview of stakeholder groups engaged with the IBD-CSIRT: Governing Bodies, Sector CSIRTS, and Constituents. Our study adopted a detailed ‘vertical’ approach via interviews and data analysis. With a ‘horizontal’ approach, we confirmed findings via interviews and a validation workshop.

relate its shortcomings to earlier work on missing feedback loops in critical security mechanisms.

### 3.3. METHODOLOGY

To answer our research questions, articulated in section 3.1, we use two approaches for this study. Using a “vertical” approach, we conduct a detailed case-study analysis of the IBD-CSIRT and its relevant stakeholders. Via a “horizontal” approach, we confirm and generalize earlier findings with other Dutch sector CSIRT practitioners. The stakeholders and approaches are presented in Figure 3.1.

Our study consists of three distinct phases, depicted in Figure 3.2. First, we analyze stakeholders engaged with the IBD-CSIRT as a case study for a detailed perspective on the provided services, expectations, and challenges.

In phase 2, we build on the results of phase 1 with an analysis of the vulnerability notification service. After the analysis, we conducted three additional interviews, two with security specialists of the Dutch National Cybersecurity Authority (NCA) notification program, and one with the IBD-CSIRT, to contextualize our findings.

Finally, in phase 3, we first interviewed other Dutch sector CSIRT practitioners. Next,

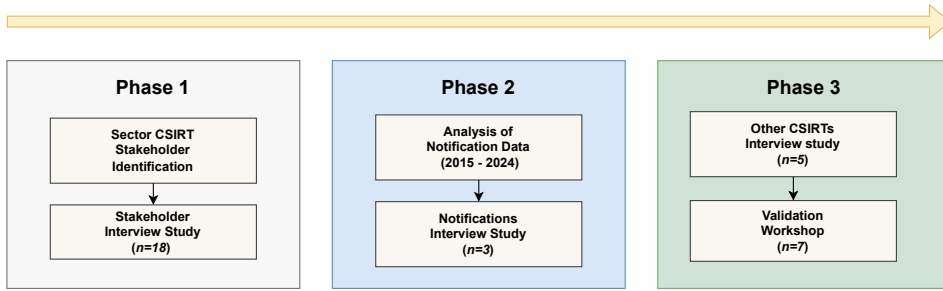


Figure 3.2: Data collection methodology.

Table 3.1: Overview of Sector CSIRTs in the Netherlands

ID	Org Name	Sector(s)	Year Founded	Org Structure	FTE 2023	# of Constituents	Budget 2023/2024
1	IBD-CERT	Municipalities	2013	Part of association of municipalities (VNG)	19	550	2+M
2	Z-CERT	Healthcare sector	2017	Private foundation	37	349	5.2M
3	SURFcert	Education and research	1992	Cooperative association	3.5	200	n/a
4	NCA	National CSIRT and sector CSIRT for government and CI	2012	National government body	280	10k+	36M
5	CERT-WM	Regional water authorities	2016	Project at Waterschapshuis	3.5	25	0.5M
6	CSIRT-DSP	Digital Service Providers	2019	National government body (Ministry of Economic Affairs)	n/a	n/a	n/a

we invited practitioners of all Dutch sector CSIRTs to a workshop session to validate our findings, which we discuss in § 3.7<sup>2</sup>.

### 3.3.1. INTERVIEW STUDIES

**Interview protocol.** We developed an interview protocol with minor variations depending on the assigned stakeholder group of the interviewee: *Governance*, *CSIRT*, or *Constituents*, provided in Table B.1. Some questions were rephrased depending on the stakeholder group. Furthermore, a minor variation was introduced when referring to either

<sup>2</sup>The Dutch National Cybersecurity Authority (NCA) acts as both national CSIRT and sector CSIRT for the sector *central government* and the sectors considered *critical infrastructures (CI)*. Under NIS2, the NCA will act as a sector CSIRT for all sectors that do not already have an established sector CSIRT.

the IBD or to another CSIRT (e.g., Z-CERT), depicted in Table B.2. In all variations, we asked participants to articulate a list of CSIRT services. The protocol contains an iteration where we ask the same questions for each articulated service.

To test our protocol, we conducted two pilot interviews, one with a participant of the IBD-CSIRT and one with a participant from a municipality. We made minor changes to the protocol after each interview. The results from the pilot interviews are not included in the final analysis. Additionally, we collected meta-information, e.g., Full-Time Equivalent (FTEs) and budget, about the sector CSIRTS. In doing so, we consulted the interview data and other sources such as organization websites and yearly reports. This information is depicted in Table 3.1.

In total, we conducted 24 semi-structured interviews with 26 security professionals between October 2024 and May 2025. Nine different municipalities were represented in the ten constituent interviews. In total, for two interviews, two participants were present. One interview with two constituent participants, and one interview with two governance participants. Two interviews were in person, and 22 interviews were conducted remotely. Each interview lasted approximately one hour. During 18 interviews, two researchers were present, while five interviews were conducted by the lead researcher alone.

**Table 3.2:** IBD-CSIRT Stakeholder Respondents ( $n = 18$ )

ID	Organization	Role	Group
P-G1	Ministry of Interior (BZK)	Policymaker	governance
P-G2	National Cybersecurity Authority	Manager	governance
P-G3	National Cybersecurity Authority	Technical Manager	governance
P-G4	National Cybersecurity Authority	Threat Data Expert	governance
P-G5	Association of Municipalities (VNG)	Manager	governance
P-S2	IBD-CSIRT (municipalities)	Security Expert	csirt
P-S3	IBD-CSIRT (municipalities)	Security Expert	csirt
P-S4	IBD-CSIRT (municipalities)	Security Expert	csirt
P-C1	Municipality	Technical Security	constituent
P-C2	Municipality	Information Security	constituent
P-C3	Municipality	CISO	constituent
P-C4	Municipality	CISO	constituent
P-C5	Municipality	Information Security	constituent
P-C6	Municipality	Information Security	constituent
P-C7	Municipality	Information Security	constituent
P-C8	Municipality	CISO	constituent
P-C9	Municipality	Security Expert	constituent
P-C10	Municipality	Security Expert	constituent

**Participant selection.** For phase 1, we identified three groups of stakeholders with whom participants interact with sector CSIRT practitioners: (1) governing bodies, (2) sector CSIRT staff, and (3) constituents. To be considered for the study, participants had to interact with the IBD-CSIRT. For CSIRT staff, we focused on practitioners handling incidents and engaged with constituents. At the IBD-CSIRT, five practitioners fitted our

**Table 3.3:** Vulnerability Notifications Respondents ( $n = 3$ )

ID	Organization	Role	Group
P-N1	National Cybersecurity Authority	Security Expert	governance (notifications)
P-N2	National Cybersecurity Authority	Security Expert	governance (notifications)
P-N3	IBD-CSIRT (municipalities)	Security Expert	csirt (notifications)

**Table 3.4:** Other (Non-IBD) Sector CSIRT Respondents ( $n = 5$ )

ID	Organization	Role	Group
P-S1	National Cybersecurity Authority	Sector CSIRTs Expert	csirt
P-S5	CERT-WM (water)	Security Expert	csirt
P-S6	Z-CERT (healthcare)	Security Expert	csirt
P-S7	Z-CERT (healthcare)	Manager	csirt
P-S8	SURFcert (higher education)	Security Expert	csirt

criteria, of which we interviewed the majority ( $n = 3$ ).

We pitched the study at a webinar for municipalities where people could opt in to be contacted. Through snowballing, we gathered additional constituent participants. Via contacts at the Dutch National Cybersecurity Authority (NCA), personal relations within our research group, and conferences, we recruited respondents from the sector CSIRT and governance bodies. In doing so, we tried to maintain a balance in the number of participants per stakeholder group. Table 3.2 describes the number of participants per group for phase 1. Note that each group has a different ID letter identifier.

For phase 2, we reached out to the NCA for participants who were involved in the vulnerability notification program. For phase 3, the NCA helped us set up a workshop session with relevant sector CSIRT practitioners. Via our network, we reached out to practitioners from other sector CSIRTs. We reached out to CSIRT-DSP for an interview, but they declined to participate because the organization is merging with the national CSIRT.

**Ethics.** A user study is our primary data source for this study. We received approval from our Institutional Review Board (IRB) for conducting this human-subjects research. We obtained informed consent from participants before conducting the interviews and the workshop. Participants were informed about the study and use of information for which they provided informed consent beforehand. We assured participants that their data was handled confidentially and would only be presented in an aggregated and anonymized form. The CSIRT community is small, but the organizations remain large enough to protect participant anonymity. No participant identities or municipality names would be mentioned in the paper. However, we do provide the participant ID with their role and organization type as context to quotes. Before publication, we presented a draft of the paper to the participants to check and correct quotes attributed to them, and to ensure no statements or data could lead to attribution.

**Interview coding.** Interviews were transcribed and coded using ATLAS.ti [12]. The inter-

views were coded inductively by two researchers using codebook-style Thematic Analysis (TA) [26]. During the entire coding process, the themes and codes were discussed periodically with the full author team to settle on central themes to ensure the reliability of findings [151].

Both researchers coded the first interview separately in the same coding session, discussing the coding afterwards to build an initial codebook. Next, the lead researcher coded nine interviews, and the other researcher coded two other interviews. Both researchers coded those interviews separately, while discussing them in periodic coding meetings. The interview coding work split was decided on the participant planning schedule and the researcher's availability. After this coding process, the researchers reviewed each other's work, discussed the results, and refined the codes. There was no major disagreement on the meaning of individual codes, but rather, codes were grouped or split for refinement. The coded interviews were then recoded with the refined codebook by the lead researcher. The lead researcher then coded the remaining ten interviews. As a quality check, the other researcher coded one of these ten interviews in isolation. Afterwards, the two researchers compared the coding of the interview. No major disagreements surfaced during the discussion. This led to another minor refinement of the codes. The coded interviews were then recoded with the updated codebook by the lead researcher. Finally, minor code refinements, i.e., code renaming, were made by both researchers during the analysis phase.

We did not calculate Inter-Rater Reliability (IRR), but did monitor the form and reasoning underpinning disagreements (as noted above); this is in line with Braun & Clarke, and others, who note that IRR is not usually a measure of quality for codebook-based TA [26, 151]. The final codebook is available in § B.1.

### 3.3.2. DATA AND METHODOLOGY FOR VULNERABILITY NOTIFICATION ANALYSIS

In phase 2 of our study, we investigate the vulnerability notification mechanism. Most vulnerability data that reaches constituents is based on data from a security non-profit called Shadowserver Foundation [242]. Shadowserver runs global scans daily and then sends reports for detected vulnerable machines to national CSIRTS worldwide. The latter are then meant to send it onwards to sector CSIRTS and other stakeholders.

We examine the notification delivery mechanism by investigating IP addresses in Shadowserver reports and comparing them to IBD-CSIRT ticketing data that captures all outgoing notifications to constituents. The IBD-CSIRT generously gave the lead researcher access to their ticketing data. These tickets do not include general security advisories but, instead, include notifications about vulnerable and compromised assets that the IBD-CSIRT receives from the national CSIRT and responsible disclosures. The lead researcher had on-premise access to the data via a VPN connection. We collected a total of 3,065 tickets, created between 17 July 2015 and 16 September 2024.

Additionally, the IBD-CSIRT gave the researchers a database with 'last-updated' changes to the asset inventories, as registered by the constituents. The asset inventory comes in two types. The first type ('ICT photo') describes the constituents' software infrastructure. The second type ("IP ranges and domains") contains the IPs and domains of the constituent.

Furthermore, the research team had access to the full Shadowserver reports for the Netherlands from 2018 to 2024. Shadowserver reports contain all IP addresses for that country for a specific vulnerability. The vulnerability, and hence report, is also typically assigned a severity level. The number of reports per severity per year is depicted in Table 3.7. Reports from Shadowserver are not static, over time, they get merged or discontinued. Since we do not consider IPv6 addresses within this study, we marked the reports associated with these addresses as “n/a”.

We received the set of IP addresses (IPv4) for Dutch municipalities from the IBD-CSIRT for 2022 and 2024. The list for 2022 contained 251,851 IP addresses for 278 unique constituents. The list for 2024 contained a total of 191,643 IP addresses for 295 unique constituents. IP lists are continuously maintained, and these are two ad-hoc snapshots that the CSIRT could share.

We matched IPs in Shadowserver reports with the IPs of municipal IPs for 2018 up to, and including, 2024. For the years 2018-2022, we used the 2022 IP list to find matches. For 2023 and 2024, we used the 2024 IP list. There is, however, a risk that IPs are outdated because IPs can be updated throughout the year. We use the asset-update database to remove results when the assets may have been updated. We checked the organization of a hit against the 'last update' field for IPs in the update database. If a hit occurred before or in the last update year for that organization, we removed it from the results. With this approach, we tried to mitigate the risk of false positives. We go into details of the findings in § 3.5.

### 3.3.3. VALIDATION WORKSHOP

To test the generalizability of our findings we organized a 90-minute focus-group session for security practitioners of all Dutch sector CSIRTs. The session included seven participants from all the Dutch sector CSIRTs we previously interviewed (Table 3.1). The lead researcher acted as the session moderator, a co-author acted as the assistant moderator. Responses were recorded via note-taking by the assistant moderator for note-based analysis [189]. Participants were asked to indicate their stance on each finding by a show of hands: (a) agreement, (b) disagreement, or (c) abstention. Following the vote, we invited participants to elaborate on their views through open-ended discussion. This approach is a form of mixed analysis [237].

## 3.4. RESULTS: CHALLENGES AND STAKEHOLDER EXPERIENCES WITH SECTOR CSIRT SERVICES

Here we address our first research question (RQ1): *What are service-specific challenges and expectations of stakeholders on the services provided by CSIRT practitioners?* We asked participants about their experiences with sector CSIRT services. We first asked what services are provided or used, resulting in a set of services articulated by participants, depicted in Table 3.5. According to practitioners, these are the valuable services that sector CSIRT practitioners provide.

Next, for each service, we asked, how is it used and what obstacles do you run into in practice? We coded the interviews and identified sub-codes for each service. Below, we iterate over the articulated services, describe how participants used them, and describe

the challenges participants ran into.

We interviewed participants from three types of stakeholder groups: *governance*, *csirt*, and *constituent*, as seen in Table 3.2. The participant identifiers, e.g., P-G2, P-S5, or P-C3, reflect which stakeholder group they belong to. The notification participants have the N identifier (e.g., P-N1).

**Table 3.5:** Sector CSIRT services derived from the interviews

Name	Description
Incident Response	Detecting and addressing security incidents to limit impact. A CSIRT may support triage, provide playbooks, assist with stakeholder management, and advise on communication.
Advisories	Non-asset-specific vulnerability information, typically including steps to patch or mitigate issues.
Expert Insights	Guidance on specific security topics, including <i>knowledge products</i> such as templates for organizational processes.
Vulnerability Notifications	Asset-specific alerts about vulnerabilities or abuse affecting constituent-reported assets.
Intelligence Sharing	Ad-hoc warnings and updates (e.g., new phishing trends), along with tips, experiences, and relevant best practices.
Outreach and Community	Activities to build relationships, foster trust, present services, and learn about constituent challenges.

### 3.4.1. INCIDENT RESPONSE: CLASHING EXPECTATIONS

All participants indicated that incident response (IR) is the most important service a sector CSIRT provides. This is in parallel to how governing bodies for local government acted on global efforts to boost cybersecurity for these sectors (P-G4). How this service functions in practice, and what constituents need, is harder to determine. While this service is deemed important by all participants, constituents were somewhat disappointed by the incident response capability that the IBD-CSIRT provided, noting that they hoped the sector CSIRT would provide off-the-shelf, ready-to-go solutions to manage and (technically) remediate the incident.

Where some constituents had not yet experienced an incident, the image they had was instead informed by a notorious case where an organization called the IBD-CSIRT for help during an incident, and the IBD-CSIRT responded to clarify that it does not send first responders on-site. P-S2 noted that constituents wrongly believed that the IBD-CSIRT has digital forensic specialists on standby during an incident. This signalled to constituents that the IBD-CSIRT, despite being perceived as an incident response party, did not come to aid during an incident in the way constituents expected. This undermined the credibility of the sector CSIRT as a trusted organization that assists constituents during an incident.

These misaligned expectations led to disappointment or angry phone calls within the constituency when help was needed. CSIRT participants noted that the reasoning was that municipalities are themselves responsible for their information security, and the CSIRT only offers additional help. The incident response capability, therefore, is not viewed the same by the CSIRT and the constituency. Different perceptions of services or

issues by different groups are not uncommon in the security domain [114, 117], nor in other domains such as healthcare [190]. While the specifics of earlier work differ, what they share is that by making perceptions and assumptions explicit, practitioners are able to overcome the challenge of misaligned expectations.

For example, CSIRT practitioners noted that they did learn from the experience, and they nowadays more explicitly communicate to their constituents (e.g., on their website, via fact sheets, during webinars, etc.) what to expect from them during an incident, i.e., triage, playbooks, advice, stakeholder support, and/or coordination efforts. This effort seems to better align the expectations, as P-C1 to P-C10 all described the IR capability in those terms. What we see here also highlights the impact expectations can have on how defenders seek to coordinate in critical situations—none of the constituent participants indicated that they expected the IBD-CSIRT to send incident responders, but this had already been informed by shared stories ahead of the CSIRT’s clarification efforts.

The IBD-CSIRT is gradually being considered to instead be more of a “trusted broker” for Digital Forensics and Incident Response (DFIR) services. As P-C6 noted, “...I expect them to refer me to a forensics company [as] they don’t have the capacity for that”, with the Sector CSIRT seemingly then learning where they can play a role in the defender community that is within their capabilities while being shaped by the needs of constituents. This is in contrast to other sector CSIRTs, where P-S5 and P-S8 indicated that their sector CSIRT *does* have the capability to send technical incident responders on-site to their constituency (due to having more advanced technical capabilities).

P-G1, P-G4, and P-S2 felt that the title “CSIRT” did not automatically require technical incident response capabilities (despite P-G2 indicating this). Their views highlight that the resources required to meet the specific needs of every constituent would be problematic, noting that the capability should be provided efficiently, meaning it can also be provided by a sector CSIRT with the help of a DFIR provider with which the organization has a contract.

### 3.4.2. ADVISORIES: CONTESTED VALUE

Similar to the incident response service, according to P-G4, historically, advisories were a primary task of a sector CSIRT, and served as an important measure to prevent, rather than mitigate, incidents. P-G1 and P-G5 noted the importance of this service, and its acknowledgement on a regulatory level: under the NIS2 directive, paragraph 3b [44], CSIRTs must disseminate information relevant to their constituency about vulnerabilities. According to paragraph 3a, the *national* CSIRT must monitor and analyze vulnerabilities. The sector CSIRTs, therefore, only need to forward those advisories to their constituents. From a governing perspective, P-G1 and P-G5 suggested that they view the distribution of advisories via CSIRTs as an important service to comply with the NIS2 directive. However, they also noted, that the legislation for the actual implementation is still under development and they are having talks with practitioners from governing bodies and CSIRTs to determine the details.

IBD-CSIRT participants noted that they are very dependent on the national CSIRT in providing advisories. They do not spend any resources on expanding on that service, such as parsing the advisory or providing new advisories. By contrast, P-S5, P-S6 from other sector CSIRTs noted that they developed their own additional advisories.

These advisories are about software products that are popular or specific to their constituency. Unsurprisingly, P-S5 and P-S6 highlighted the value of the advisories to their constituency. By contrast, P-S2 questioned the value of the advisories of the national CSIRT, stating “...we don't want this service to be a day job on our end, nor send too much information to the constituents...we want to provide value... it's a waste of our time if they can't act on the advisory or if they get the information already from elsewhere.” Similarly, P-S8 stated that much of the information from CSIRT advisories can be found elsewhere, and therefore, his sector CSIRT stopped sending advisories altogether.

The contested value of advisories is observed among municipal constituents, too. Some constituents perceived the advisories as one of the critical services of the sector CSIRT, as it helps them learn about potentially vulnerable systems in their network. Other constituents mentioned that they get the same information, sometimes sooner, from other sources. Particularly, participants of larger municipalities considered this a flaw of the service, or even of the sector CSIRT itself. They would rather receive advisories directly from the national CSIRT than from the sector CSIRT because they would get it faster.

According to constituent participants, several other factors influence the perceived value of advisories. First, P-C1, P-C3, and P-C6 reported that there is a process in the organization set up to directly turn advisories into tickets. These tickets are used for reporting and to allocate staff to act on the advisory. P-C6 noted “...we use them for our internal processes. When a notification comes in, a ticket is created. We can act on those tickets, and I can report on them by the end of the month.” While this benefit may seem mundane, getting such patching processes in place is challenging for many constituents, according to P-C1.

Second, because the advisory is sent from an authoritative organization, it helps convince internal stakeholders, even when the issues have already been flagged by the internal team. For example, P-C2 noted, “sometimes they don't take our word for it, and then the authority of the CSIRT is very useful.” Advisories sent from CSIRTs thus carry authoritative value for their constituency. The value of institutional authority does not appear to be a Dutch phenomenon, as it was also observed in [11], where CERT-SE was widely regarded as a trusted source of intelligence. Furthermore, it appears that the advisories have a utility beyond the actual contents of the advisory: practitioners acknowledge the advisory as a valuable tool to convince internal stakeholders. At the constituency, the advisory serves as a common ground between technical practitioners and management. A similar dynamic was found in [276], where the authors found that the National Vulnerability Database (NVD) worked best as a boundary object for a dialogue between the security team and others within the organization.

### 3.4.3. EXPERT INSIGHTS: ONE-SIZE-FITS-NONE

Participants from governing bodies and the CSIRTs noted that they have deep expertise in both security and the constituency. They offer this security expertise to their constituency in the form of topical briefs, webinars, policy documents, templates and advice, threat assessments, and quarterly news updates. Within this domain, practitioners often also use the term “knowledge products”. Contrary to the incident response and advisory services, this service is not part of a regulatory framework but grew organically.

According to P-S2, this service originates from the CSIRT observing low levels of security knowledge among constituents, a relatively high turnover rate of municipal staff, and municipalities with many more responsibilities than resources. P-S3 also noted that the CSIRT wanted to grow professionally and in its service offerings, reason from the perspective of the needs of its customers. In doing so, they ran customer satisfaction surveys to determine what customers (i.e., their constituency) needed. A need for these "knowledge products" was one of the findings from those surveys, and therefore got allocated more resources by the CSIRT.

However, CSIRT participants explained that there are operational challenges for the sector CSIRT in providing knowledge products to its constituency. P-S2 and P-S3 stated that their sector CSIRT offers a lot of value to their constituents with "templates" for policy and internal processes, such as a standardized supplier agreement or the outline of a security process. However, they need to provide these templates to many constituents, with many maturity levels. This makes it problematic to provide specific, actionable templates that anyone can use. The result is that, according to constituents, the templates are often generic and impractical. For example, P-C5 stated, "*...the information could be very useful but...it's almost impossible to provide this 'golden glove' for all organizations...*". Producing these types of products has another inherent drawback for sector CSIRTs, it burdens them with the maintenance of the documents. These one-size-fits-all solutions are not uncommon in the security world. In [147], the authors found that one-size-fits-all solutions to provide more security are not realistic, because it places burdens on their users, which leads to great variations among participants' security approaches and implementations. This also has parallels to policymakers' views of consumer advice around the security of smart devices [256], in wanting to be seen to provide advice, yet it does not match the needs of all the people who need it.

#### 3.4.4. VULNERABILITY NOTIFICATIONS: THE UNSEEN SERVICE

CSIRT participants and constituents mentioned that the sector CSIRT sends asset-specific vulnerability and abuse notifications to its constituents. P-S2, P-N1, and P-N2 explained how this process works. This process, i.e., "notification flow", is visualized in Figure 3.3 and further detailed in section 3.5.

Interestingly, constituents rarely mentioned this service during the interviews, suggesting that they do not see this as a valuable sector CSIRT service, whereas CSIRT participants highlighted it as a service to prevent incidents. When the service did come up, participants indicated that they "probably signed up" for the service but had not received any notifications. Some constituents assumed that because they did not receive notifications, things were probably okay.

Constituents may receive vulnerability notifications from other sources (e.g., vulnerability scanners), reducing the perceived value of CSIRT notifications. Although we did not ask participants directly about alternatives, some noted that they conduct their own vulnerability scanning, which could lessen their reliance on the CSIRT. Still, none described the CSIRT's notifications as redundant. Instead, three constituents (P-C1, P-C2, P-C7) viewed them as a useful "second opinion" or as complementary to their external attack-surface management. This suggests that, when executed well, CSIRT vulnerability notifications retain clear value for constituents. This also suggests that CSIRT notifi-

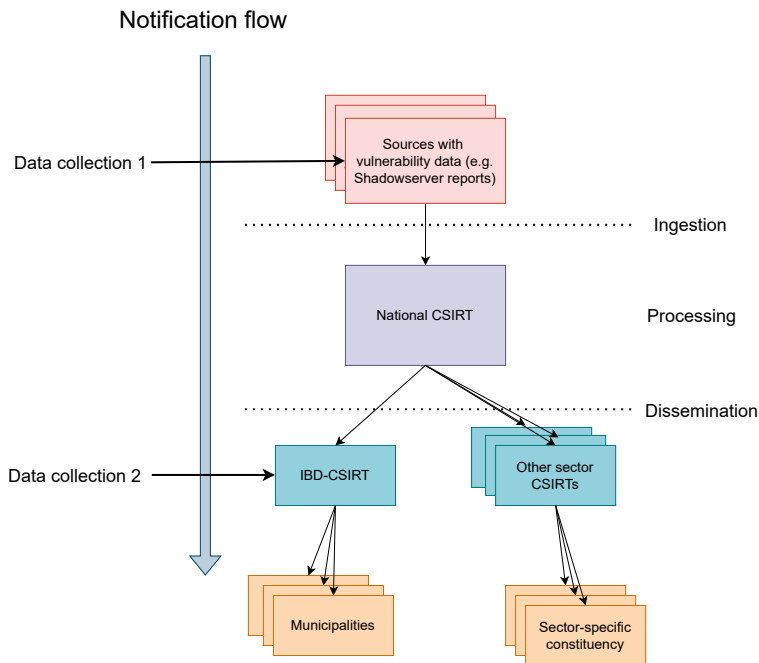



Figure 3.3: Overview of automated vulnerability notification flow

cations have value for highly resourced countries, or sectors, e.g., banks, which could be explored in future work.

For most constituents, the service did not really seem to be on their radar. Somewhat related, P-S2 noted that there were times when the sector CSIRT rarely received notifications from the national CSIRT, which made them wonder if everything was ok.

To send notifications, CSIRTs need an up-to-date asset inventory of constituents. We hypothesized that if constituents valued the service, they would maintain a recent asset inventory. Therefore, we asked participants of the IBD-CSIRT for any data about constituent asset registrations. They maintained a list of the last asset change per constituent in Excel, as described in section 3.3. We found that this data shows that constituents rarely update their IPs and domains with the sector CSIRT, depicted in Figure 3.4. There are approximately 350 municipalities, Figure 3.4 shows that more than 100 municipalities last updated their IPs at the sector CSIRT in 2014. Another explanation is that constituents lack the resources to track all their assets. This is indeed the case—many reported limited capacity for various security tasks. In [14], it was also observed that defensive security efforts were hindered by budget constraints. However, all constituent participants maintained an asset inventory, even if it was potentially incomplete or inaccurate, which was not updated at the CSIRT.

This data suggests there might be a bootstrapping problem on the notification service: if the constituents do not keep their asset registrations current, then they might not get the relevant vulnerability and abuse notifications; yet, in the absence of getting



**Figure 3.4:** Constituency last-update value for asset registration per year. To illustrate, in 2014 more than 100 municipalities (out of approximately 350 total) last updated their IP ranges at the IBD-CSIRT.

relevant notifications, they are not incentivized to keep the registration current. At the same time, though, it’s also the constituents that have the least (technical) knowledge and capacity and may need help.

Lastly, CSIRT participants explained that vulnerability notifications can also come in via responsible disclosures, as opposed to notifications from the national CSIRT. Several CSIRT participants mentioned the value of their responsible disclosure service, protecting their constituency from the hassle of managing responsible disclosures themselves. They described the sector CSIRT process as follows. The sector CSIRT receives responsible disclosures from researchers who think the issue is related to some municipal asset – e.g., it relates to a domain name with the name of a municipality. The sector CSIRT then verifies the claim and assesses the severity. It then sends an actionable report to the constituent. Several CSIRT participants noted how much time and energy this service costs the sector CSIRT. The cost resides mainly in verifying the vulnerability and discussing the severity with the reporter. P-S6 noted that 90% of the received vulnerabilities were not taken into consideration after verification.

While CSIRT participants took pride in the handling of responsible disclosures, the

service was not mentioned by constituents. This underlines that the perceived value of a sector CSIRT service, at least partly, depends on the visibility of the service by the constituent – i.e., most constituents have not received a responsible disclosure notification in recent years.

### 3.4.5. INTELLIGENCE SHARING: VISIBILITY EQUALS AUTHORITY

CSIRT participants explained that advisories focus on vulnerabilities that are not asset-specific. Vulnerability notifications are asset-specific warnings. Expertise insights are often strategic and technical organizational products. By contrast, they noted that *intelligence sharing* revolves around the sector CSIRT sending out ad-hoc warnings for an issue (e.g., a new type of phishing attack), tips, experiences, best practices, or other relevant information.

A majority of constituent participants noted that they frequently received such intelligence and considered this sector CSIRT service very valuable, helping them tune their defenses, share expertise, or get help. P-S4 noted that, despite the service being valued, it was sometimes hard to get information back from the community. He suspected that constituents are afraid that they might do things wrong and that their information and ways of working are judged. In fact, he stated, the CSIRT lauds any efforts from constituents that give back to the community. In [278], it was observed that malware analysts also tended not to share information. However, in their case, they refrained from sharing because they felt it was only one-way or that it did not help others.

This service is highly visible—constituents frequently receive intelligence—which may explain its perceived value. An alternative explanation is that utility, rather than visibility, drives value perception. Indeed, utility is likely to be a factor. However, as discussed in subsection 3.4.4, responsible disclosure handling was considered useful by CSIRT practitioners but was not visible to constituents and therefore not perceived as valuable. This supports the hypothesis that visibility is also a factor in how constituents value a service. According to participants from the CSIRT and governing bodies, sharing intelligence has another benefit for the sector CSIRT: the sector CSIRT continues to maintain its reputation as an expert or authority within its domain or sector.

### 3.4.6. OUTREACH AND COMMUNITY: BRINGING CONSTITUENTS TOGETHER

CSIRT participants stressed the importance of knowing their constituency. They mentioned that they travel across the country to attend meetings, join regional initiatives, meet people from constituent organizations, and talk to suppliers. P-S3 and P-S4 noted that underpinning these efforts is a desire of the sector CSIRT to foster trust with and among constituents, showcase services, and learn what issues constituents are dealing with. These efforts are part of the ‘constituent expertise’ that participants from governing bodies value so highly of sector CSIRTS.

Constituents considered sector CSIRT outreach and community-building efforts more practically, focusing on how the sector CSIRT has a facilitating role in bringing constituents together who face similar challenges. Some constituents stated that regional initiatives failed without the sector CSIRT involved, and noted that the sector CSIRT should have a facilitating role. P-S4 noted that the sector CSIRT would like to personally know all constituents, but that simply is not possible with the size of the constituency

and their geographic distribution over the country. Regional initiatives, like a CISO platform where local CISOs get together, help the sector CSIRT scale their networking efforts as well as bring constituents together.

Another CSIRT participant highlighted the value of these efforts to manage expectations among constituents. With the NIS2 directive coming into effect, many more organizations will fall under a critical sector. These organizations, generally, do not have experience with cybersecurity legislation and obligations. However, those organizations will have obligations. P-S1 noted that “...we want to make ourselves as clear as possible and show what services we provide...I want that clearly on the website...I want to manage those expectations, and clearly communicate what we do. Tell them that they are a [responsible] ‘entity’, not a customer.” This illustrates the underlying relation between sector CSIRT practitioners and constituents: CSIRT practitioners are there to support constituents as best as they can but ultimately, constituents have their own responsibilities, and they are not CSIRT customers who can demand whatever they like.

**Table 3.6:** Overview of the Value and Challenges for Articulated CSIRT Services

Service	Value	Challenges
Incident Response	Assistance during an incident CSIRT acts as trusted (DFIR) broker	Misaligned expectations over service implementation Role of CSIRT during incident unclear Lack of resources
Advisories	Incident prevention NIS2 compliance Advisories facilitate constituent internal processes	Dependency on national CSIRT Advisories not timely Custom advisories cost scarce resources Constituents cannot always act on advisory Advisory information available elsewhere
Expert Insights	Sector-specific expertise available Provide and preserve cybersecurity expertise	Uniform knowledge products not actionable Service maintenance costs resources
Vulnerability Notifications	Incident prevention Asset-specific notifications	Dependency on national CSIRT Constituents do not update assets Service not valued by constituents Responsible disclosures efforts not visible and valued
Intelligence Sharing	Sector-specific actionable intelligence Establishes CSIRT as an authority	Intelligence collection from constituency
Outreach and Community	Foster personal relations between CSIRT, constituency and suppliers Develop expertise among constituency	CSIRTs role in regional initiatives unclear Constituency too big and dispersed CSIRT can only facilitate

### 3.5. RESULTS: EVALUATING THE VULNERABILITY NOTIFICATION SERVICE

In subsection 3.4.4, we found that CSIRTs and constituents did not appreciate the vulnerability notification service equally: CSIRTs considered it important, in stark contrast

to constituents who were mostly unaware of it. To understand this dichotomy, we investigate how this service functioned in practice. We do so by measuring notifications at each stage of the flow ( Figure 3.3) to quantify throughput and identify obstacles. We cover the period 2015 to 2024. In late 2024, the system was revamped. Our methods were: (1) a longitudinal empirical analysis to assess whether notifications reached IBD-CSIRT constituents and which factors impeded delivery; and (2) three follow-up interviews to contextualize findings: two with national-CSIRT program leads and one with a sector-CSIRT practitioner operating the service.

3

**3.5.1. MEASURING NOTIFICATION FLOW**

The dissemination of vulnerability notifications from the sector CSIRT to constituents happens via a ticketing system at the IBD-CSIRT. We assess the functioning of this flow by comparing the notifications that constituents *should* have received to what they *actually* received.

The methodological details of this analysis are discussed in § 3.3.2. We analyse ticketing data from IBD-CSIRT. The system creates a ticket for all vulnerability notifications received from the national CSIRT. It then automatically notifies the associated municipality of the ticket. We analyzed 2,826 tickets ranging from 17 July 2015 to 16 September 2024.

Most vulnerability data received from the national CSIRT came from Shadowserver. For our evaluation, we had access to full Shadowserver vulnerability reports that the national CSIRT receives, the asset inventory list of constituents, and the IBD-CISRT ticketing data. This allows us to check if IP addresses that show up in Shadowserver reports (data collection point 1) are seen in tickets of the IBD-CSIRT to their constituents (data collection point 2 in Figure 3.3).

First, we analyzed the vulnerability data from Shadowserver. It consists of reports that contain IP addresses that are vulnerable or compromised [242]. Table 3.7 depicts the total number of reports by year and severity. The national CSIRT receives all Shadowserver reports for the Netherlands. The researchers had access to the same Shadowserver data. According to P-N3, the IBD-CSIRT does *not* have access and relies on the national CSIRT.

**Table 3.7:** Number of Shadowserver Reports

Severity	Year						
	2018	2019	2020	2021	2022	2023	2024
<b>info</b>	2	2	3	6	9	9	10
<b>low</b>	5	5	6	4	4	6	6
<b>medium</b>	12	14	16	19	24	24	23
<b>high</b>	22	24	24	28	40	41	43
<b>critical</b>	9	9	9	22	19	18	22
<b>special</b>	0	0	0	1	4	1	0
<b>n/a</b>	4	4	2	11	24	36	35
<b>total</b>	<b>54</b>	<b>58</b>	<b>60</b>	<b>93</b>	<b>124</b>	<b>135</b>	<b>139</b>

Using the IPs in the Shadowserver reports, we determine the set of ‘hits’: municipal IP addresses that show up in the Shadowserver reports. These hits make up the set of notifications that IBD-CSIRT should have received.

According to the national CSIRT participants P-N1 and P-N2, a ticket is created once per day per report type. When a vulnerability notification ticket is created, it contains an abuse attachment that may contain multiple IPs. This leads to only one ticket per day. Therefore, we take tickets per day as the unit of analysis. We analyzed 2,826 tickets; of these, 378 describe notifications pertaining to the Shadowserver data.

We matched all Shadowserver reports on municipal IPs and found hits in 67 unique reports. However, IBD-CSIRT received notifications for only 6 different reports. Thus, our first important finding is that many Shadowserver reports do not lead at all to notifications to the constituents.

For the six report types that do show up in the tickets, we observe that there are hits with municipal IPs that do *not* lead to ticket creation. From 2018 to 2024, we observed 378 tickets, i.e., 378 days where a ticket was made that references one or more hits on municipal IPs in the Shadowserver data. In the raw Shadowserver data, for the same period, we observed 1,365 days for which Shadowserver registered a hit on one or more municipal IPs in the six reports. Thus, only 27% of days with Shadowserver hits led to ticket creation. Figure 3.5 depicts the number of days per year that a hit was present in the Shadowserver data, but no ticket was created. This observation is most salient in 2022, when 353 days with one or more hits did not lead to any ticket creation. Figure 3.5 depicts the number of notifications that should have been sent versus the number of notifications in the ticketing data.

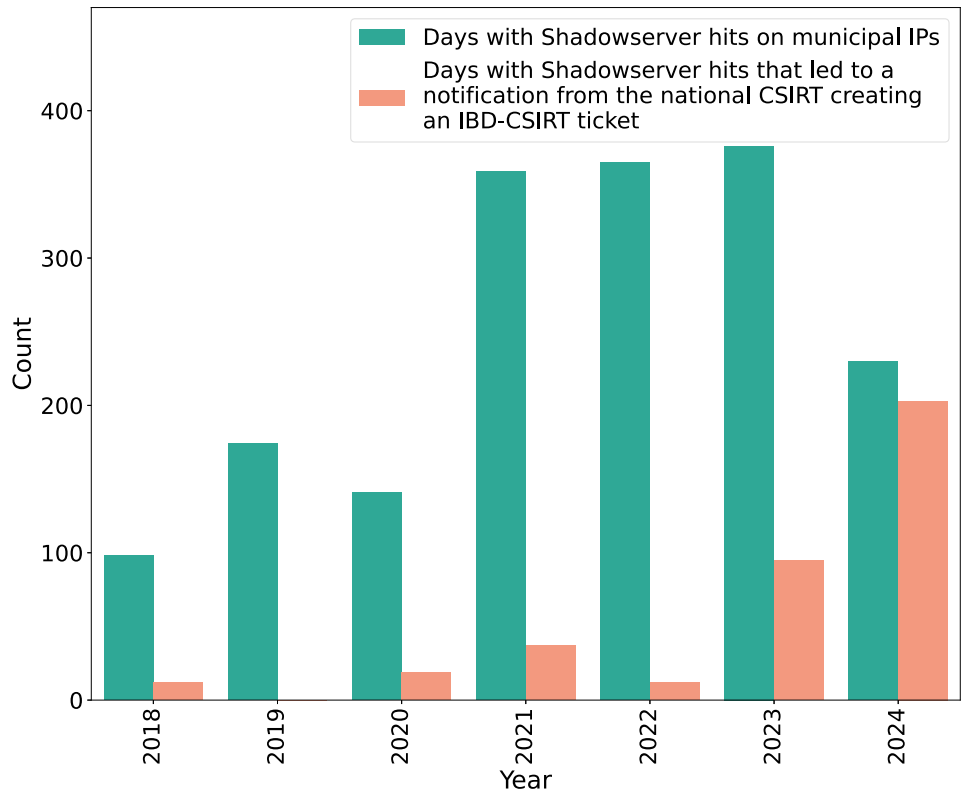
These findings suggest that the vulnerability notification service was not functioning as intended: most vulnerability data is not reaching the respective constituents. This may explain why constituents rarely mentioned the service during the interviews: they probably never received notifications. Surprisingly, though, this mechanism was in place for many years during which the missing notifications were not detected.

### 3.5.2. REFLECTIONS EVALUATING THE NOTIFICATION PROGRAM

We found that a lot of vulnerability notifications were not reaching constituents, contrary to the prescribed notification flow depicted in Figure 3.3. We presented these findings to P-N1 and P-N2 of the national CSIRT, who operated the program. We also presented the findings to P-N3 of the sector CSIRT. We wanted to confirm we correctly understood the notification mechanism and identify the root causes for the missing notifications.

P-N1 and P-N2 talked openly about the challenges in the setup and running of the notification service. P-N1 indicated that, contrary to other domains, the cyberdomain was new at the time, and nothing in terms of regulation or best practices was in place. They were tasked to do something new and had to figure out what worked along the way.

For our first finding, that only a fraction of the Shadowserver reports were sent, they noted that, depending on the type of organization to which they send the notifications, the system may decide to include more or fewer reports. This system was put in place to not overburden recipient organizations like IBD-CSIRT by sending too many notifications that they could not act on. N3 was not aware of this filtering process, let alone consented to it.



**Figure 3.5:** All days with Shadowserver hits on municipal IPs vs. the days with Shadowserver hits on municipal hits that led to an IBD-CSIRT ticket.

For our second finding, not all vulnerable IPs led to a ticket, P-N1 and P-N2 stated that they did not directly have an explanation for the numbers. They speculated that their notification tooling used e-mail to send the Shadowserver reports to IBD-CSIRT. Perhaps some tickets were not created due to delivery failures. They also noted that their internal tooling sometimes has issues and may temporarily be offline. It could well be that a “queue” of messages could have been deleted. Additionally, they mentioned that their internal notification tooling was set up to process hits without automatically sending actual notifications, as a safety precaution.

When discussing the total lack of notifications in 2019, the participants speculated that for a certain period, an IBD-CSIRT recipient address to send notifications to may have been missing or “unchecked” in the tooling. They could verify this checking/unchecking mechanism in the tooling but could not verify whether this had taken place. They also noted that one of the six Shadowserver reports was only “turned on” in 2023, explaining some of the discrepancies. To verify, we removed this report from our analysis for the years 2018-2022 and observed that some missing notifications in 2021 and 2022 indeed decreased. However, many missing notifications remain, in those and in other years.

Finally, P-N1 and P-N2 noted the lack of quality assurance processes to test and evaluate their notification pipeline. While they voiced a desire for such processes, a lack of resources and more pressing priorities prevented their implementation. The absence of a feedback loop is therefore problematic.

The absence of a feedback loop in the critical mechanism of vulnerability notifications has been a fundamental problem that has been observed across many countries. Earlier studies on such mechanisms [34, 133, 257, 280] found that recipients almost never respond to notifications, so it is very difficult for the sender to observe if the notifications were correctly delivered, let alone acted upon. The lack of a feedback loop has also been observed as a problem in government efforts. When CISA evaluated its Ransomware Vulnerability Warning Pilot (RVWP), it observed that its personnel relies on subsequent internet scans to infer that the vulnerability was mitigated, confirming the absence of a feedback loop from the notified entity [251].

In sum, the practitioners who run a national notification service are faced with internal decision-making processes (i.e., filtering hits from reports by making assumptions about the recipients), source selection and processing (what reports are trustworthy), and tooling issues. These are factors that impede notification delivery. Practitioners are faced with limited resources and diverging priorities. The lack of feedback caused the missing notifications to go undetected for years. Even if the sector CSIRT and its constituents have strong bonds, they do not know they are not receiving the information meant for them, so they can't correct the problem.

### 3.6. RESULTS: STRATEGIC CHALLENGES FOR SECTOR CSIRT PRACTITIONERS

In this section, we address our second research question (RQ2): *What are the strategic challenges for practitioners of a sector CSIRT in providing services?* We interview practitioners and determine the challenges that transcend specific CSIRT services.

Two groups of organizational challenges emerged from grouping subcodes from the interviews, available in Table B.3: (i) governance structure and external stakeholders; and (ii) infrastructure and capability management.

#### 3.6.1. GOVERNANCE STRUCTURE AND STAKEHOLDERS

We observe that most sector CSIRTs originated from practitioners' best-effort bottom-up (i.e., unregulated) initiatives with limited resources, providing many services. Additionally, P-G4 noted that the sector CSIRTs have a minimal top-down structure and rely heavily on personal relations. Consequently, existing sector CSIRTs have different governance structures, depicted in Table 3.1. By contrast, the NCA did not originate as a bottom-up organization.

The sector CSIRTs' governance structure can be problematic. For example, Z-CERT (CSIRT for healthcare), under NIS2, will be entitled to government funding. However, because of its legal structure, a foundation, the organization may not be eligible to receive the funding. It may also not recommend certain commercial parties to constituents, which is considered 'disrupting the market'. This highlights the importance of considering the risks and benefits of various governance structures when establishing the sector

CSIRT because it may limit the specific things a sector CSIRT may undertake down the line.

For certain services, the sector CSIRT is highly dependent on other stakeholders, e.g., for advisories described in § 3.4.2. Constituents depend on those advisories as a mitigation measure to manage their attack surface. For some constituents, according to P-C5 and P-C6, those advisories are the primary source of vulnerability information. A similar dependency exists between the sector CSIRT and the national CSIRT in providing vulnerability notifications. Finally, constituents rely on commercial parties for the IT infrastructure. This creates a situation wherein a constituent may receive advisories, but is incapable of acting on them because they do not directly control the infrastructure. This process is problematic in practice, as the IT service provider does not always act as fast as the organization wants. For example, P-S2 reported that organizations sometimes have to wait up to a month before their ticket on a patch is resolved.

Next, we observed trust as an important component during an ongoing incident. Previous academic work also identified the essential role of trust among CSIRTS [8, 13, 88, 124, 165, 272], and in industry [181]. CSIRT participants noted that if an organization is hit by an attack, the sector CSIRT is supposed to be invited to a triage meeting. However, CSIRT participants noted that this does not always happen for a variety of reasons. For example, an organization may be scared because the sector CSIRT acts as an independent party, and victim organizations fear that they might be blamed for not having security measures in place.

Closely related, CSIRT participants noted that they want to be a *trusted partner* for their constituents. *Liaisons* were mentioned as a way to build trust. Additionally, sector CSIRTS engage in various outreach and community efforts, described in § 3.4.6. Participants from sector CSIRTS fear that they are viewed as an authoritative organization that tells constituents what to do. Therefore, CSIRT participants note that they take great care to protect trust among stakeholders. Yet, constituents did not indicate that they feel the sector CSIRT is telling them what to do too much. Instead, some constituents indicated the sector CSIRT may be *more* authoritative to get basic security measures in order at certain lagging constituents.

Regarding its constituent population, a sector CSIRT faces challenges in serving many constituents with mixed maturity levels. All CSIRT participants acknowledged this, even those with fewer constituents, see Table 3.1. The value of CSIRT services can suffer because organizations have varying needs—knowledge, tools, or infrastructure. For example, P-C4 and P-C5 noted that the IBD-CSIRT uses “one-size-fits-all” knowledge products. This approach fails both immature organizations, which lack the expertise to use them, and mature ones, for whom the products are too basic. P-C8 observed that the products may work for smaller organizations but not for larger ones. Or, as P-C4 stated, “*one size fits none*”.

Finally, P-G1 emphasized that the diversity of constituents is where a sector CSIRT adds value: acting as a glue, building trust, and facilitating collaboration—a view shared by all CSIRT participants.

### 3.6.2. INFRASTRUCTURE AND CAPABILITY MANAGEMENT

Participants state that infrastructure and asset management is challenging. Maintaining a complete, up-to-date asset inventory is difficult. While internal IP ranges are stable and monitored, SaaS services frequently change, and inventories rely on departments reporting assets. These unreported systems weaken security and hinder sector CSIRTs. Without accurate inventories, CSIRTs cannot deliver or receive vulnerability notifications from the national CSIRT. Many services are outsourced or moved to the cloud, and constituents rely on external security services. P-S8 warned that this reduces organizational control and visibility, making it hard to act on alerts. P-G5 noted that during incident triage, debates arise over control, outsourcing, and responsibility for security measures.

**Table 3.8:** Overview of CSIRT Strategic Challenges

CSIRT Strategic Organizational Challenges
Mandate and legal structure may impact service provision and resource allocation
Limited resources for many services
CSIRT is dependent on other organizations
Lack of trust may impede service provision
Diverse constituency with mixed-maturity levels
Outsourcing parts of constituent infrastructure makes incident triage complex

## 3.7. RESULTS: FINDINGS VALIDATION WORKSHOP

To more rigorously test the generalizability of our findings and further deepen them, we validated findings with practitioners from multiple sector CSIRTs via a workshop session. The practitioners were all security officers. In light of the time constraints, we selected a subset of seven findings of § 3.4 and § 3.6, which we paraphrased for brevity and easy comprehension. We used the approach detailed in subsection 3.3.3.

An overview of the results is depicted in Table 3.9. In many cases, some participants voted "abstained" because they felt the agree / disagree vote was too simple. Instead, they preferred to elaborate on their situation.

**1. Practitioners of sector CSIRTs are dependent on other stakeholders for the delivery of services.** Two participants noted that this finding depends on the specific service, mentioning their red-team services as an example where they are not dependent. Sector CSIRT practitioners may thus vary in their level of dependency on others depending on the specific service. As a "coordinating" body [77], however, by its very nature, many services (e.g., notifications or advisories) introduce these dependencies.

**2. The diversity of constituents is a big challenge in the development and offering of services.** Those who agreed raised their hands instantly and vocally agreed with this statement. The participant who disagreed claimed that his organization simply offered several "flavors" of its services, depending on the maturity of the constituent. Some of

**Table 3.9:** Validation Workshop Voting Results per Finding

Finding	Section	Agree	Disagree	Abstained
Practitioners of sector CSIRTS are dependent on other stakeholders for the delivery of services.	6.2	4	0	3
The diversity of constituents is a big challenge in the development and offering of services.	6.2	5	1	1
Trust between stakeholders is essential for practitioners of sector CSIRTS to provide services.	6.2	3	1	3
Many constituents rarely update their asset list at the sector CSIRT which undermines the effectiveness of sector CSIRT services.	6.3	2	2	3
Constituents expect technical assistance on location with the Incident Response service, but practitioners of the sector CSIRT are not equipped to do so.	4.1	3	1	3
Practitioners of sector CSIRTS had concerns about not receiving all vulnerability notifications of the national CSIRT.	4.4	1	2	4
Services are often shallow because scarce resources are distributed among many services.	6.2	5	0	2

the participants who agreed responded that providing such “flavors” of services took additional resources that are rarely available.

**3. Trust between stakeholders is essential for practitioners of sector CSIRTS to provide services.** The participant who disagreed argued that practitioners from his organization can *always* deliver services, but without trust, the quality of those services would be degraded. However, he did not elaborate on how.

All participants agreed that trust in the security domain is important because much information is shared via personal connections for information sharing. Furthermore, sector CSIRT practitioners put great value in trust with their constituency, so they are contacted by them when needed, so that the CSIRT can be ‘found’ by constituents, as noted in sector CSIRT industry guidelines [180].

Thus, personal relations are important for unhindered information exchange between stakeholders. For those relations, trust is essential, and it acts as a catalyst in providing information-dependent services.

**4. Many constituents rarely update their asset list at the sector CSIRT which undermines the effectiveness of sector CSIRT services.** A discussion occurred over what constitutes an asset. Some participants considered ‘contact details’ assets, and they had trouble with outdated details and could not reach constituents. Others noted that they invest time in managing those contacts, and they did not recognize the problem. A participant who disagreed indicated that their organization controls the network infrastructure of its constituents, so IP registrations were not much of a problem.

If we expand the term ‘assets’ to include contact details, then asset registration is widely seen by participants as problematic because they are not frequently updated. Inaccurate contact information in the notification process, which prevents contacting an entity, was also observed by [34]. Keeping asset registrations up to date was difficult for everyone, except for the sector CSIRT practitioners who operate network infrastructure for its constituents.

**5. Constituents expect technical assistance on location with the Incident Response**

**service, but practitioners of the sector CSIRT are not equipped to do so.** One participant noted that they are equipped for IR, but it depends on the magnitude of the incident. Another participant mentioned that their documentation on how much support constituents could expect from the sector CSIRT during an incident was the biggest source of confusion among constituents [168]. The discussion focused on constituent incident response expectations of the CSIRT. This was largely seen as problematic because the practitioners could not deliver the expected support.

**6. Practitioners of sector CSIRTs had concerns about not receiving all vulnerability notifications of the national CSIRT.** Two participants noted that their organization had their own Shadowserver feed and do not rely on the national CSIRT for vulnerability data. Another participant noted that his sector CSIRT, up to now, has had no reason to doubt that they are not receiving all vulnerability notifications. This finding seemed specific to IBD-CSIRT practitioners, as only they observed anomalies in their ticketing data in 2019. However, without some kind of feedback mechanism, it is very difficult for practitioners to actually know that they are not receiving all notifications, unless there are glaring anomalies.

**7. Services are often shallow because scarce resources are distributed among many services.** Participants vocally agreed to recognizing this finding. The two participants who abstained may also have nodded agreement, but the researchers are not sure. Nonetheless, this issue was widely agreed upon.

### 3.8. DISCUSSION

In this study, we try to lift the veil on how the professionals in and around sector CSIRTs experienced their functioning. For our first research question, we analyzed the experiences and challenges tied to six specific services that sector CSIRTs provide. The results are summarized in Table 3.6. For our second research question, we identified strategic challenges that transcended the specific services, summarized in Table 3.8. Here, we want to make sense of the underlying dynamics that shape these experiences and challenges. We organize these dynamics around three concepts: resources, legitimacy, and dependency.

First, we found misaligned expectations and challenges in providing the services that are associated with the label ‘CSIRT.’ Sector CSIRTs are small, much smaller than a ‘normal’ CSIRT, because they were formed bottom-up and are funded primarily by organizations in the sector pooling some resources. This puts persistent constraints on their capabilities. While they provide various services, many of them are rather shallow in their implementation. The clearest example is that constituents expected “boots on the ground” during incident response. The sector CSIRT was not able to provide this.

A seemingly straightforward strategy to better align the expectations with the available resources would be to make more explicit to constituents what the sector CSIRT can and cannot do. To some extent, this is what happened for the “boots on the ground” issue. Over time, constituents learned not to expect this from their sector CSIRT. Such a strategy, however, would overlook the other dynamics at work, which help to understand why the misalignments not only arise, but are sustained over time.

A second dynamic is the sector CSIRT’s pursuit of legitimacy. Telling their constituents to expect less also makes the sector CSIRT less relevant. So the professionals operate in

a tension: promise too much and risk disappointment, versus promise too little and risk being irrelevant. This tension makes an easy alignment of expectations with capabilities difficult. There will always be pressure to do more for their constituents, stretching the limited resources.

Yet, it is not all about servicing the constituents. As one sector CSIRT professional phrased it: *“I want to manage those expectations, and clearly communicate what we do. Tell them that they are a [responsible] ‘entity’, not a customer.”* If the constituents are not customers, then what are they? The term ‘entities’ is a nod to the regulatory frameworks, most notably NIS2, which applies to ‘essential’ and ‘important entities.’ To call the constituents responsible entities is to say that it is their responsibility to meet certain security requirements. The sector CSIRT is meant to support them in bringing about this outcome. In other words, it is not just about meeting the needs of the constituents, but also about getting them to change their practices; paternalistically nudging them in the ‘right’ direction.

This is the reason why governments in many countries have encouraged setting up sector CSIRTS. In the EU, under NIS2, this is now even mandated. The organizations are seen as instruments to improve security in the sector. And yet, when we consider this as a community of practitioners, it aligns with commentary on community-owned interventions [23], that community-level transformation would fare better if the people within that community can own the change themselves, rather than have it forced upon them by outside experts. The sector CSIRT must understand them, show them, and act alongside them.

This dual mandate – helping constituents while also getting them to change their behavior – means that the legitimacy of the sector CSIRT comes from the top as much as from the bottom. It helps us understand why the sector CSIRT sees the vulnerability notification service as very important, yet it was barely mentioned by the constituents as something they expect from the sector CSIRT. This service is not driven by demand from the constituents, but by demand from the government. So the sector CSIRT professionals are managing two legitimacy relationships at the same time. This is also why misalignments are not easy to resolve, because alignment with one side might cause misalignment with the other. As noted by Kocksch et al. [115] in the discussion of care in IT security, “IT security, as an organizational achievement, relies on an intricate entanglement of care and organizational authority”; we get the sense that the sector CSIRT cannot walk away from their remit, this being that constituents must be protected.

Finally, we observed dynamics around dependencies. The malfunctioning vulnerability notification service clearly show a dependency on the national CSIRT. It did not forward the relevant notifications to the sector CSIRT, which in turn sent fewer notifications to constituents. This contributed to the latter not really seeing the value of the service. It also eroded the incentives for the constituents to update the asset registrations that they have submitted to the sector CSIRT.

These dependencies interact to create a bootstrapping problem: without relevant vulnerability notifications, the constituents are less likely to correctly register their assets, and the absence of up-to-date asset registration makes it less likely that they will get relevant vulnerability notifications. The more general version of this bootstrapping problem is: the sector CSIRT can provide value for the constituents if the constituents

invest time and effort into working with it, yet as long as they do not see the value, they are unlikely to make those investments.

This assumes that constituents actually know their assets to begin with. Our participants noted that constituents' users may spin up assets without telling anyone. In examining UNICEF logistics, Jack & Jackson [108] refer to the reality of 'messy infrastructures' rather than something neater. To force neatness into asset management would be to force a change to the way whole organizations operate, undoing their naturally messy nature.

In the end, the sector CSIRTs can only offer different types of carrots, they lack any kind of stick. Within the EU, the NIS2 legislation might help overcome the bootstrapping problem as it requires constituents to take certain actions – e.g., by reporting incidents to the sector CSIRTs. It remains to be seen if these obligations are enough, as some constituents simply lack the knowledge or resources to engage more fully, or to impose their own mandate upon the rest of their organization.

**Recommendations.** Our recommendations are organized around the concepts of resources, legitimacy, and dependency.

- *Focus on least-capable.* CSIRTs have limited resources, yet face a diverse constituency in terms of maturity. Given that their goal is to augment the capabilities of the constituents, it makes sense to spend the limited resources on the least-capable, rather than trying to support everyone. Such an approach leads to “one-size-fits-none” solutions. At the same time, realistic baselines should be established for those constituents to avoid raising expectations that the sector CSIRT cannot live up to.
- *Include challenges in guidelines.* Second, we recommend to extend industry guidelines by flagging the issues that our empirical data has surfaced around resources, legitimacy, and dependencies, which might help new CSIRTs to better diagnose the issues they are facing. The current support is very sparse and limited to a single SEI report about how to set up a sector CSIRT, not how to operate one [180]. The FIRST framework has yet to release guidelines specifically for sector CSIRTs [76]. The SEI report anticipates challenges and suggests that sector CSIRTs should focus on offering a small set of services and doing those well. We found that, in reality, due to legitimacy tensions, CSIRTs might overexpand their provided services, undermining the trust of constituents. Such lessons learned may help other organizations to navigate difficult issues, rather than having to reinvent the wheel.
- *Build feedback loops.* We found that practitioners could not easily detect the missing vulnerability notifications: there was no feedback mechanism. In our case, this allowed a malfunctioning pipeline from the national CSIRT to the sector CSIRT to persist for years. The absence of a feedback loop has also been observed in notification research [34, 133, 257, 280]. This work found that the final recipients often do not act on vulnerability notifications. Both problems undermine the effectiveness of this important sector CSIRT service.

This lack of a feedback loop likely applies to numerous other provided services. By introducing a feedback mechanism, stakeholders will a) know that services are

operating correctly (e.g., all vulnerability notifications are arriving), and b) know if the service is either useful or being acted on (e.g., CSIRT will know if constituents are doing something with advisories and vulnerability notifications).

- *Mix top-down with bottom-up incentives.* Fourth, CSIRTS face a bootstrapping problem. This issue can be addressed via a mix of top-down and bottom-up incentives. One top-down incentive is to make updated asset registrations a stronger norm. The Dutch national CSIRT is working towards a constituent registration platform for all sector CSIRTS and constituents [167]. This national platform is inspired by NIS2 and will replace the ad-hoc efforts of single sector CSIRT which, at least for IBD-CSIRT, were not very successful. A stronger incentive is the option to impose regulatory requirements on constituents to engage with the sector and national CSIRT, as is the case under NIS2 where constituents are required to report incidents. Bottom-up, as a community of practitioners, sector CSIRTS may continue to invest time in helping constituents develop their capabilities that will allow them to see the value of the services and encourage them to invest in engaging more with the sector CSIRT.

**Limitations.** Our research design faces several limitations, most notably regarding external and internal validity. First, our main focus was on one sector CSIRT in one country, with its constituency and governance bodies. Generalizability to other sectors in the country is supported by our cross-sector validation workshop, where practitioners from other domains reported similar dependencies, asset-inventory pain points, and expectation gaps.

We did not research sector CSIRTS in other countries. While this limits the generalizability of our findings, we would argue that the findings have wider relevance. Institutional contexts for sector CSIRTS differ across countries, yet there are strong similarities. The sector CSIRT model is literally a model, one that has been copied worldwide – first under the name sector CERT (Computer Emergency Response Team), but after the name CERT was trademarked, the label sector CSIRT was adopted. The U.S. Department of State commissioned a report by SEI to support this dissemination by developing a general framework for the founding of new sector CSIRTS [180]. Within FIRST, there is a vibrant community that supports setting up new instances of the model in different countries and sectors. Of course, all implementations will be adapted to local conditions, but we would argue that the triad we surfaced — resources, legitimacy, and dependencies — offers a transferable analytic frame that supports the work of sector CSIRT professionals elsewhere.

In terms of internal validity, our sample size was limited. We were able, however, to recruit representatives of all selected stakeholders, including all sector CSIRTS except one. Our participants only reflect a part of the perspectives within their respective organizations. We addressed this limitation by recruiting several people from an organization whenever possible. In the first phase, we interviewed three IBD-CSIRT practitioners. While a small number, it reflects the team's limited size: only five practitioners handle incidents and interact directly with constituents, with the remainder in support or administrative roles. Thus, our sample covers a substantial share of the relevant professionals. Furthermore, given the challenges of recruiting participants, limited resources,

and saturation of themes, we believe the sample size is adequate to support our findings. **Future work.** We presented exploratory work on practitioner challenges in operating a sector CSIRT. We propose several areas for future work. First, this study could be replicated in other countries or sectors to corroborate our findings or determine factors that affect our observed challenges. Second, we investigated the vulnerability notification service in detail. Yet, for the other services, an in-depth study remains. Third, for the small set of Shadowserver reports that *were* used to send notifications, the program did improve in 2024. Consequently, this may impact service perceptions by constituents, which may be analyzed further.

**Conclusions.** We investigated challenges of practitioners at sector CSIRTs by asking what stakeholders expected. In doing so, we identified challenges practitioners face in providing services. Sector CSIRT practitioners need to deal with a diverse constituent population, trust issues, and organizational dependencies. For the vulnerability notification service, this dependency turned out to be problematic, as not all notifications arrived at constituents. We highlighted factors that undermine the national notification mechanism. While regulatory frameworks increasingly rely on sector CSIRTs, there is a need to better understand these organizations as institutional structures to mitigate cyber threats.



# 4

## APT TO DISAGREE: A COMPARATIVE ANALYSIS OF ATTRIBUTION IN COMMERCIAL TI

*Attributed cyber threat intelligence (TI) plays an important role in the effective mitigation of cyber attacks. Yet, despite the central role of attribution in policy, practice, and vendor reporting, little is known about the coverage and reliability of attribution by threat intelligence vendors. No study has systematically investigated attribution across a large set of leading TI vendors. We close this gap and provide a longitudinal comparative analysis across 13.5 million IOCs collected over the last 14 years from seven vendors. To compare IOC attribution across vendors, we normalize heterogeneous feeds and reconcile actor names using an evaluated and augmented version of MISP Threat Actor Galaxy (MISP TAG). Next, we address two questions: (i) what is the scope of actor-tracking by vendors, and (ii) how consistent is attribution among vendors? We find that the majority of actors tracked by one vendor are not tracked by the other. Furthermore, IOCs observed by multiple TI vendors are rare (1%), illustrating that commercial TI feeds, like open-source feeds, primarily provide singleton IOCs. We also find limited overlap in IOCs for jointly tracked actors by two vendors. We measure attribution agreement among vendors with Krippendorff's  $\alpha$ . We find mostly moderate agreement among vendors for actor attribution. By contrast, country attribution has high agreement. Our results have implications for actor-centric defenses, compliance, and geopolitical uses of attribution.*

---

This chapter has been published as: **Ethembaoglu, A.M.**, van Wegberg, R.S. & Zhauniarovich, Y. & van Eeten, M.J.G. (2026). "APT to Disagree: A Comparative Analysis of Attribution in Commercial TI". In *Proceedings of the IEEE Symposium on Security and Privacy (S&P '26)*.

## 4.1. INTRODUCTION

Attributed Indicators of Compromise (IOCs) are pieces of forensic data that indicate a system or network may have been breached by a particular threat actor. They play an important role in high-end threat intelligence (TI), enabling defenders to detect, respond to, and anticipate threats more effectively, while also informing strategic and legal decisions. Unlike non-attributed counterparts, attributed IOCs provide organizations with insights into who is behind an attack and why it is taking place [68, 109, 222, 247].

Attribution of threats to actors is critical to various use cases. First, best practices advocate that organizations adopt TI to know the adversaries targeting them, so that they can focus their resources on tailored defences [48, 146, 182, 197, 224], thereby directing resource expenditures and improving defense effectiveness. This threat-led approach has also been incorporated into government and corporate policy guidance [15, 247, 249]. Second, attributed IOCs are used to assess legal exposure and sanction checks [254]. Some specific threat actors might be sanctioned by policy makers, such as the U.S. Office of Foreign Assets Control [78], and enterprises have to adapt their engagement (including ransom payments) and reporting obligations accordingly. Third, attribution has geopolitical implications [9]. Tying attacks to certain states or organizations has real-world consequences such as embargos on countries, export controls, sanctions, or exerting diplomatic pressure [172, 226].

Established models – Pyramid of Pain, the Q-model, and MITRE ATT&CK – differentiate short-lived, easily changed low-level IOCs (e.g., IPs, URLs) from more persistent high-level behavioral patterns, i.e., Techniques, Tools and Procedures (TTPs), typically treating the latter as more reliable for attribution [53, 159, 209]. However, recent studies show TTPs alone rarely discriminate actors, and practitioners often ground investigations in low-level indicators [106, 131, 209, 217, 254]. Consequently, low-level IOCs remain indispensable to contemporary attribution and the decisions derived from it.

Despite the central role of attribution in policy, practice and vendor reporting, little is known about the coverage and reliability of attribution by threat intelligence vendors. Disagreements in attribution risk undermining the very ‘threat-led’ defenses they are intended to support. While there is a wealth of prior work on threat intelligence, focusing on various aspects of quality like accuracy, timeliness, and coverage [126, 152, 243], attribution is rarely mentioned in this research and, to the best of our knowledge, has never been evaluated. This is likely related to the fact that the bulk of this work is based on open source TI, which lacks the attribution effort of high-end commercial vendors. Only one study [25] has analyzed the products of two leading threat intelligence vendors. It briefly touched on attribution and found little overlap in the IOCs of the two vendors when tracking the same actor. However, to date, no study has systematically investigated attribution across a large set of leading TI vendors.

In this paper, we aim to close this gap. We conduct a longitudinal comparative analysis of 13.5 million unique, attributed IOCs collected over the last 14 years (2011-2025) from seven commercial TI feeds. Determining the quality of the attribution of IOCs is severely impeded by the lack of ground truth data on attackers that hampers all research on TI [152]. In addition to absent ground truth, the attribution process itself is confidential, using data and knowledge that vendors do not make public. That makes independent evaluations very difficult. Yet, what we can do is analyze and compare the

attribution claims that different vendors make about their IOCs. Simply put, if these claims conflict, then at least one of those attributions is incorrect. We seek to answer two research questions – RQ1: *what is the scope of actor-tracking by vendors?* and RQ2: *how much agreement is there in attribution across vendors?*

We find that the majority of actors are tracked by only one vendor. When vendors do track the same actor, they report sets of IOCs with very low overlap. When multiple vendors report the same IOC, the level of agreement in their attribution is poor – as indicated by Krippendorff’s  $\alpha$ , a metric for inter-rater reliability. Agreement levels improve to satisfactory when we remove temporary attribution labels. The agreement about which country is behind the attack is consistently high among vendors. In sum, we make the following contributions:

- We present the first systematic investigation of attribution in high-end commercial threat intelligence feeds, providing a longitudinal comparative analysis across 13.5 million IOCs of seven vendors.
- Recent work found that the lack of a shared industry mapping was hampering analysts [216]. We validate and augment the leading industry actor mapping, MISP Threat Actor Galaxy (MISP TAG or TAG), which contains 855 actors. While not exhaustive, our validation shows TAG has high accuracy when checked against a deconflicted ground-truth mapping from CrowdStrike and Microsoft and against vendor TI reports for a random sample of 50 actor names.
- We find that vendor coverage of threat actors is limited. Each vendor covers only a fraction of the actors (2-17%) in TAG. The union of all seven feeds, which would cost USD 1-2 million per year, is able to track just 34% of the actors from TAG.
- We find that only 1% of all attributed IOCs are observed by more than one vendor. Even when vendors claim to track the same threat actors, the IOCs they report for those actors show minimal overlap. When vendors do observe the same IOC, their attribution agreement is moderate: Krippendorff’s  $\alpha$  ranges from 0.76-0.81 for two vendors, even when the observations are just a week apart in time. Agreement on country attribution is consistently high, with  $\alpha$  ranging from 0.92-.95 for two vendors. There is poor agreement on IPv4 indicators ( $\alpha$  ranges from 0.65-0.71), but even on MD5 hashes, disagreement remains.
- We identify implications for security practices around the use cases of attributed TI.

## 4.2. RELATED WORK

Our work resides at the nexus of studies on threat intelligence, attribution, and Advanced Persistent Threats (APTs).

**Threat Intelligence.** There is a rich line of work on, primarily, open-source TI. Early efforts focused on abuse feeds and block lists and found limitations in accuracy, timeliness, and coverage [126, 152, 243]. More recently, studies have described TI IOCs from open and commercial sources and measured them along various quality dimensions, including, among others, coverage, timeliness, relevance, overlap, delay, volume, and accuracy [25, 95, 118, 136, 220, 261]. Studies that measure the quality of TI all suffer from the lack of ground truth [152]. The authors of [221] conducted a qualitative study to identify the most effective ways to present TI quality to analysts. Bouwman et al. [25] conducted a mixed-method study to determine how customers value commercial TI. They

found that TI wasn't exclusively used for detection purposes but also to understand the threat landscape. To our knowledge, it is the only study analysing IOCs of reports from two leading commercial vendors. TI promises to enable the detection of attackers within defenders' networks. It assumes that once an actor is observed in one network, it may be observed elsewhere [107]. Several studies examined the effectiveness of IOCs for timely intrusion detection and tried to determine the optimal lifespan of IOCs [5, 22, 116, 246]. The authors of [22] found that IOC lifetimes were not static and developed a technique to dynamically estimate lifetimes.

**Frameworks and models.** Several models include the use of TI for cyber defense. The 'pyramid-of-pain' model illustrates that some indicators are harder to replace for attackers than others. Low-level IOCs can be quickly replaced, and TTPs are therefore believed to be the hardest to change. The Q-model of Thomas Rid relates to attribution and identifies three different levels: tactical, operational, and strategic. On the tactical, and most technical, level, Rid observes that low-level IOCs play a key role at the start of attribution investigations [209]. In [226], Steffens describes the fine-grained MICTIC framework for technical attribution. It is similar to the Diamond model, a framework to analyze cyber intrusions [28].

**Attribution.** Some prior work has focused on tools for (APT) attribution [148, 214, 268]. The authors of [203] proposed CSKG4APT, a knowledge graph model for attribution using Open-Source Cyber Threat Intelligence (OSCTI). In [212, 266], the authors applied machine learning to classify APT groups using features from sandbox reports and string and code features, respectively. Scrutinizer identified unknown APT samples through code reuse analysis [157]. In another study, the authors proposed ADAPT, which utilized static features extracted from heterogeneous file types for attribution by clustering executables and documents in threat groups and campaigns [215]. In [266], the authors developed explainable APT attribution for malware using NLP techniques.

There is another strand of research about attribution as a practice done in industry and in the context of international relations [16, 90, 138, 209]. More recently, [216] shed light on the perceptions of security practitioners in analyzing APT-level threats, including attribution. Other attribution-related work focuses on empirical evaluations of various claims within industry. Recent studies [217, 254] showed that the use of TTPs to distinguish between threat actors has its limitations. In [93], the authors found that a big challenge in attribution is the unreliable naming and labeling of threat actors and their associated TTP. Leite et al. demonstrated that name server patterns may be useful for clustering attacks [131]. All these studies find that low-level IOCs remain essential for attribution, though some point to a mix of low and high-level IOCs [277, 279]. Indeed, respondents of a SANS survey indicated that low-level indicators of compromise were valued higher than information about TTPs [106].

**Attacker Economics and APTs.** Various studies looked into the operations of attributed attackers. In [128, 245], the authors found that most attack campaigns described in TI use known vulnerabilities. This supports the idea of [6] that, in general, attackers operate cost-efficiently. By contrast, the authors of [18] found that APT-attributed malware used more sophisticated evasion techniques.

In this work, we close research gaps in existing work by conducting an empirical and longitudinal study of seven feeds of commercial TI. Previously, only [25] studied

two commercial TI feeds. However, no study has systematically investigated attribution across a large set of leading TI vendors.

### 4.3. BACKGROUND

As threat actors carry out attacks, they leave traces within victim systems and networks. Security firms collect victim logs and analyze malicious artifacts to identify clusters and patterns. The attribution of attacks to specific actors is a dynamic, ongoing process, as vendors continually acquire new data and insights. These may originate from additional victim systems, other customer networks, collaborations with vendors, or publicly available sources such as forums and the dark web.

In [226], Steffens explains the role of ‘intrusion sets’ in the analysis and attribution of cyber attacks. Malicious activity can be analyzed by observing both low-level indicators and TTPs being used together repeatedly. Such co-occurrences are termed intrusion sets. One actor can be associated with more than one intrusion set – i.e., it can use different approaches for an attack. For example, a phishing e-mail attack versus a watering hole attack. Each attack has its own artefacts, malware, control servers, etc.

Security firms collect a catalog of intrusion sets against which they can compare new attacks. Historical data is of high value for correlating incidents [216]. If TTPs and/or indicators match or overlap, the probability that the same group is behind the attack increases. An intrusion set is typically assumed to refer to a single actor, even if the specific identity of the threat actor may still be unknown. This process explains why threat actor attribution is so dynamic and has a high degree of uncertainty. There is no ground-truth about attackers, only telemetry data and probabilities [226]. For this reason, some vendors provide a type of ‘work in progress’ name or label, for intrusion sets that are not yet attributed to a specific named actor. Customers can use the “early attributed” IOCs for block lists, detection signatures, and patching priorities [259]. Vendors each use their own convention for these ‘work in progress’ (WIP) names. For example, Mandiant uses UNC[XXXX] (e.g., UNC1878), Recorded Future and Microsoft use TAG-[XX] and DEV-[XXXX], and Trend Micro uses VOID[X] [194, 259].

Vendors all use their own naming schemes for actors. For example, with the SolarWinds supply-chain intrusion in 2020, Mandiant tracked the actor as UNC2452, Microsoft as NOBELIUM, while CrowdStrike labelled the associated actor APT29 [160]. Security firms have their own customer base and therefore observe attacks that might or might not have been observed by other firms as well. This also means that some of their telemetry data is unique. Analysts at different security companies may focus on different aspects of a set of attacks, resulting in differently delineated intrusion sets. For this reason, the vendor-specific naming scheme, at least initially when uncertainty is high, makes sense.

Yet, vendors do not report on actors in a vacuum; they collaborate, and their analysts read and refer to other vendors’ TI reports [216]; thus, gaining insights from each other’s reporting. Consequently, in their own reports, they often provide the reference names (aliases) used by other vendors for that actor. This information can be used to create a mapping that links different vendor-provided names referring to the same actor.

Such mappings are also maintained by the community and by reputable vendors, including MISP Threat Actor Galaxy (TAG) [57]; APTMap (which uses TAG and other

sources) [47]; Florian Roth's spreadsheet [213]; Thai CERT [32]; MITRE [159]; Breach-HQ (769 actors) [104]; Qianxin [202]; and vendor-specific mappings [156, 223]). However, all mappings suffer from a lack of ground-truth. They present someone's best understanding that two different names used by different vendors represent the same actor, while keeping up with a dynamic threat-actor landscape. Earlier work, in the context of AV, highlighted difficulties in measurement of AV performance when vendors provide different virus labels [143, 161]. Nonetheless, these mappings are widely used in academic research [25, 127, 144, 211, 217] and within industry [145, 171].

## 4.4. METHODOLOGY

To perform this study, we partnered with a national government enterprise organization with several thousand employees that purchases TI feeds for monitoring and analysis. The partner agreed to collaborate under two conditions: (i) anonymity of the partner organization and (ii) anonymity of the commercial TI sources they acquired. We discuss more details, implications and ethical considerations regarding this collaboration in section 2.3. All feeds are sourced by the partner organization from top-tier commercial threat intelligence vendors that focus on APT tracking. The partner organization subscribes to multiple top-tier commercial feeds to increase coverage. Vendor selection is driven by enterprise procurement decisions, not research cherry-picking. We used IOC data that was available on the internal systems of the partner organization, no IOCs were filtered by the research team.

We examine actor and country attribution of IOCs published by seven commercial TI vendors (designated as V1-V7), including only the records that contain an explicit actor attribution by the vendor. Table 4.1 shows the number of attributed IOCs per vendor. In total, our dataset comprises 13,685,963 actor-attributed IOCs spanning 14 years. However, the collection windows of IOCs per vendor vary, as described in Figure 4.1. For all vendors except V6, the last IOC was collected in June 2025. Note that for six vendors, we have IOCs spanning for at least five years or longer. This overlap in time means they were making attributions in the same time frame.

While we are not allowed to disclose the vendor names, we can state that five of the seven vendors appear in the MISP list of threat intelligence producers [158]; three are listed as NCSC-UK Assured Cyber Incident Response Providers [248]; and four are covered by Gartner Peer Insights for Threat Intelligence Products and Services [89]. Each vendor appears in at least one of these sources. Vendor V7 has a regional focus, whereas the others do not. Within industry, these feeds are considered top-tier, as is reflected in their pricing. The annual fees per feed (including reports) vary between \$150,000 and \$350,000.


### 4.4.1. NORMALIZATION OF HETEROGENEOUS FEEDS

Vendor feeds arrive in heterogeneous formats, including flat JSON-like lists, STIX 1.x/2.x objects [17], and MISP Event/Object data [265]. We transform each feed into a flat schema (Table 4.2) to enable cross-vendor comparison.

We first normalized their IOC data and put them into a single scheme. We renamed any indicator type referring to an IPv4 address to IPV4, any domain to DOMAIN, and any

**Table 4.1:** Overview of actor-attributed IOCs per vendor. Vendor actor coverage is relative to TAG's total of 855 actor objects. Parentheses indicate WIP labels.

ID	Total IOCs	Unique Actors	Actors in TAG	Coverage of TAG	IOCs with Actor in TAG
V1	382,991	721 (599)	147	17%	80,361 (21%)
V2	12,633,525	229 (0)	146	17%	12,599,531 (99%)
V3	85,089	110 (0)	88	10%	54,915 (64%)
V4	384,109	105 (14)	68	8%	64,723 (16%)
V5	11,335	131 (0)	38	4%	2,616 (23%)
V6	22,312	130 (0)	89	10%	17,701 (79%)
V7	12,110	39 (0)	18	2%	7,009 (57%)



**Figure 4.1:** Overview of collection window timelines

URL to URL. Next, the feeds were inconsistent in what hashing they applied: some vendors included multiple types (e.g., SHA1, SHA256, MD5) within a single IOC record, while others created separate records for each type. Additionally, vendors used different types of hashes. However, all vendors provided at least an MD5 hash in most cases. Therefore, we use the MD5 hash type to compare among vendors. We omitted 8,130 IOCs (0.06% of all attributed IOCs) that did not provide an MD5 hash type.

The preceding steps resulted in four indicator types (IPV4, DOMAIN, URL, MD5) that were present in every feed. We refer to those types as the *primary indicator types*. In our dataset, 13,423,585 indicators belonged to the four primary indicator types (98%). Figure C.3 depicts the number of IOCs per indicator type per vendor. For the remaining 2%,

**Table 4.2:** Normalized IOC schema.

Name	Description
id	Universally unique IOC identifier
vendor_id	IOC identifier from vendor
created	IOC creation date (UTC)
type	Indicator type (e.g., domain, ipv4)
value	Indicator value
vendor	Anonymized vendor identifier
actor	Vendor attributed actor name
actor_id	Normalized actor ID
country	Attributed country (optional)

we identified 55 other (secondary) indicator types. This included, for example, a YARA rule or an X509 certificate. Two feeds included combinations of primary indicator types, such as the type `domain|ip`. We converted an IOC record with such an indicator type into two IOC records, one with the indicator type set to `DOMAIN` and one with the indicator type set to `IPV4`, copying the other fields. We converted 6 combined-type IOCs to primary indicator types. We did not apply any transformation to the other indicator types.

#### 4.4.2. ACTOR NAME MAPPING

Within this study, we use the MISP TAG mapping, specifically a snapshot from 29 July 2025 with commit 42b5d56. We use this source for several reasons. First, MISP TAG is an open project that is actively maintained [92]. Second, it provides information on the largest number of actors (Table C.1 in section C.1). Third, TAG is recognized and used by other popular tools and platforms. For example, Malpedia, a vetted curation platform offering a corpus of labeled, unpacked malware samples, uses TAG in mapping actors to malware [145].

However, our preliminary analysis indicates that this mapping still has room for improvement. Therefore, below we describe the steps taken to validate TAG and enhance the mapping using TI vendor reports.

##### ACTOR OBJECTS, NAMES AND WIP LABELS

TAG uses the following JSON-like structure to describe an actor [92]:

```
[{"name": "<actor_name>",
  "synonyms": ["<alias_1>", "<alias_2>", "..."],
  "country": "<country>"
}, ...]
```

Note that `name` in TAG serves as both a name and also as the identifier. To make an explicit distinction between actor objects and actor names (the set of all associated names, including aliases/synonyms), we posit that an actor object has an ID and a list of names, which are all the names, synonyms, and aliases for the actor. Therefore, for our analyses, we adopt a simplified format to avoid ambiguity between the `name` and `synonyms` fields that exist in TAG:

```
[{"<id>": {
  "names": [<all_known_actor_name_and_aliases>],
  "country": <country_or_null>
}, ...]
```

In this schema, `<id>` is an auto-incremented identifier for the actor object; `names` is a deduplicated set of strings with names, aliases, and synonyms for that actor. In this study, when we refer to an *actor*, we imply an *actor object*.

Last, as described in section 4.3, various vendors attribute IOCs with a type of "work in progress" or "uncategorized" label, which we dub a WIP label. Usually, an actor with a WIP label gets merged with a known actor. For example, this is the case with UNC1945, which later was linked to actor LightBasin. In reality, though, not all WIP labels get merged. In TAG, in 29 cases (out of 855, 3%), an actor has a WIP label and no synonyms. In our study, if an actor occurs in only one feed, has a WIP label, and has no synonyms,

we refer to it as a WIP actor. We identify this subset so we can later look at attribution agreement with and without this set of actors.

#### COUNTRY ATTRIBUTION

One vendor provides an explicit country attribution in the feed. In six out of seven feeds, a country label was not explicitly a part of the IOC feed. For four vendors, country attribution can be inferred from actor naming conventions (e.g., "Bear" for Russia, "Panda" for China). Two remaining feeds did not provide country attribution explicitly or through naming conventions. For these, we used the country of the actor as mapped in the TAG dataset; otherwise, we derive it from the corroborating vendor reports used to extend the mapping. Because not all actors are state-linked, the country can be null.

#### TAG VALIDATION AND AUGMENTATION

Given the importance of TAG for our analysis, as well as its value for practitioners in industry, we validate and augment the TAG data. Specifically, we check for ambiguous mappings of names to actors and compare the TAG mapping against explicit claims from the vendors themselves. We also augment TAG by adding additional names and mappings from vendor reports. To achieve this goal, at first, for each actor not mapped in TAG, we looked up the latest 3 to 5 TI reports of the vendors that mentioned that actor. Next, for those actors, we identify aliases from the TI reports. For example, a report might say "ActorX AKA ActorY", or "ActorX (ActorY, ActorZ)", or "ActorY, which we track as ActorX", or variations hereof. These examples are not exhaustive, and we encountered many small variations. For this particular example, these statements would allow us to link the names "ActorX, ActorY, and ActorZ" to the same actor with multiple names.

#### 4.4.3. AGREEMENT METRICS AND ANALYSIS

Our goal is to quantify cross-vendor agreement on actor attribution for the same IOCs. We define co-observation of an IOC across vendors as an exact match on the normalized *indicator value* field – e.g., `evildomain.com` or `192.0.2.34`.

However, we should also take a time dimension into account because IOCs have a limited lifespan. If an IOC is observed once by one vendor and once by another, but those observations are far apart in time, then they are less likely to have been part of the same threat. So when those vendors attribute them to different actors, that might not signal disagreement about the attribution but simply mean that the threats belong to different actors. For this reason, we want to look at agreement for co-observed IOCs where the observations by different vendors are temporally close to each other.

How close? There is prior work on IOC lifetimes [5, 22, 116, 246], but no consensus on estimates of IOC lifetime. Therefore, we analyze our data to see how different lifetimes affect the agreement levels. Given the fleeting nature of IOCs, we introduce several observation periods (OPs): 7, 14, 30, and 60 days. For example, for OP=7, we analyze IOCs that were observed by two or more vendors within a period of 7 days. This way, we can see if the agreement level is sensitive to the temporal distance between the observations of the same IOC.

For the co-observed IOCs, we compute inter-rater reliability (IRR) statistics on nominal labels to determine the level of agreement among observers. We treat the attribution

of co-observed IOCs as a classification task, where each vendor assigns an actor label to an IOC. We measure agreement using Krippendorff's alpha [122], which accommodates multiple raters and nominal data. It takes chance agreement into account, penalizing any kind of disagreement. Krippendorff's alpha ranges from -1 (perfect systematic disagreement) to 1 (perfect agreement). To qualify the level of agreement, we interpret Krippendorff's alpha as it is conventionally understood in the context of IRR [149]:  $\alpha < 0$  indicates systematic disagreement,  $\alpha = 0$  indicates chance-level agreement,  $0 < \alpha < 0.67$  indicates poor agreement,  $0.67 \leq \alpha \leq 0.79$  indicates moderate agreement,  $\alpha \geq 0.80$  indicates satisfactory agreement, and  $\alpha = 1$  indicates perfect agreement. We compute IRR statistics for sets of two, three, or four vendors co-observing IOCs, as no IOCs are co-observed by five or more vendors. We estimate 95% confidence intervals (CI) for  $\alpha$  using a nonparametric bootstrap with 500 resamples, to determine the reliability of the  $\alpha$  values.

4

Furthermore, to determine the agreement among specific vendors, rather than among any pair of vendors, we also calculate pairwise agreement. For each vendor pair and period, we include IOCs for which both vendors assigned an actor label, count matches versus mismatches (i.e., assign the same actor or country), and report the proportion of matches. This measure is a percent agreement in  $[0, 1]$ , not chance-corrected (unlike Krippendorff's  $\alpha$ ), where 1 indicates perfect agreement.

Finally, in subsection 4.7.2, to get a better understanding of observed disagreement, we use VirusTotal (VT) [263] and look up 867 hashes to determine the *file type* and *malware label*.

## 4.5. TAG VALIDATION AND AUGMENTATION

An accurate mapping of vendor actor names is vital to compare attributed IOCs. So as we use TAG in this study, we validate its quality by i) checking TAG for ambiguity in actor names, and ii) measuring the validity of the mapping against claims of the vendors themselves.

### 4.5.1. ACTOR AMBIGUITY IN TAG

TAG defines mappings between actor objects and their associated names (aliases or synonyms). In total, TAG contains 2,080 distinct names linked to 855 actor objects. However, these mappings are not always one-to-one: some names are associated with multiple actor objects, introducing ambiguity in the name-to-actor relationship.

For example, the name Evasive Panda appears both as the primary name of one actor object and as an alias for another actor, BRONZE HIGHLAND. As a result, Evasive Panda cannot be treated as a unique identifier for a single actor. A naïve approach might be to merge actor objects that share a name, assuming they refer to the same entity. Yet, such merges can propagate inconsistencies, as, once merged, the combined alias set may overlap with additional actor objects, causing cascading ambiguities that obscure the underlying actor structure.

We refer to this phenomenon as “alias ambiguity” (AA). In total, from the 2,080 actor names in TAG, we found 23 such ambiguous names (1%), presented in Table C.4 in section C.1. We excluded the associated 22,477 IOCs (0.2% of all IOCs).

### 4.5.2. VALIDATING TAG

We want to test whether TAG accurately reflects the claims of the vendors themselves when it comes to mapping different vendor names to the same actor. In other words, do the vendors concur that these mappings are accurately reflecting their understanding? We perform two validation steps to check the accuracy of TAG in representing the vendor's own statements about a shared attribution. First, in June 2025, CrowdStrike and Microsoft announced a bilateral reconciliation of actor names, publishing a cross-vendor mapping (CS/MS Mapping) [154]. We compare TAG against this mapping. Second, we randomly sample 50 actor names from TAG and check TAG mappings against the claims of the vendors in their reports.

#### CROWDSTRIKE AND MICROSOFT MAPPING

The CS/MS Mapping published by CrowdStrike (CS) and Microsoft (MS) lists 84 actor mappings, where a name used by CrowdStrike is matched against the corresponding name of Microsoft. We matched each CS/MS name pair to TAG and asked whether both names resolve to the same TAG actor – either as the primary TAG name or within the synonyms. Overall, 83% (70/84) of pairs were present and consistent with TAG. Four pairs, so eight distinct names, were absent from TAG (i.e., item A in Table C.5). We added these to our augmented TAG.

Eight pairs were present in TAG but resolved to different TAG actors (i.e., item B in Table C.5), i.e., the CS/MS Mapping treats each pair as names of a single actor, whereas TAG treats them as names of two distinct actors.

For example, in TAG, Ruby Sleet is described as an actor linked to North Korea's Ministry of State Security, and Kimsuky as a North Korean actor that targets South Korean think tanks, industry, nuclear power operators, and the Ministry of Unification for espionage purposes. These actors could potentially be the same actor or actors in different divisions within the same organization, the North Korean Ministry of State Security.

We consider the second case (item B) similar to the “actor ambiguity” described in subsection 4.5.1, and therefore exclude IOCs with those actor names. This affects 103,815 IOCs (0.9% of all IOCs).

#### VALIDATING AGAINST TI REPORTS

To verify the quality of TAG mappings, we randomly sampled 50 actors and their names from TAG. For each sampled actor and its names (aliases), we examined the latest TI reports that mentioned the aliases to extract explicit alias statements. Such statements are, for example, when a vendor name is followed by “AKA” and a list with names of other vendors. This indicates that the vendor itself sees these as the same actor. Another example is the phrase “we track X as Y”. We compared all such statements with the TAG entries. We observed no discrepancies. In all 50 cases, the mappings that the reports explicitly contained, were all consistent with TAG.

### 4.5.3. TAG AUGMENTATION

We examined how many of the actors from a vendor feed are listed in TAG. The number of TAG-listed actors per vendor is depicted in Table 4.1. For example, 147 (of all 721 actors tracked by V1) are listed in TAG, with 80,361 (21% of all attributed IOCs by this vendor) IOCs attributed to those actors.

The results show that all vendors have a portion of their actors not showing up in TAG. Thus, we aimed to augment TAG. We do so by examining vendor TI reports for known aliases for those actors that occur in the vendor feeds but are not listed in TAG, using the format described in subsection 4.4.2. The result of our augmentation is a new mapping that consists of the original TAG mappings, plus additional actors and aliases.

#### AMBIGUITY IN MAPPING

In our augmentation process, we ran into ambiguous mappings done by vendors. For example, a vendor stating ‘ActorX (AKA ActorY, ActorZ, ActorC)’ may be problematic when ActorY and ActorZ are associated with distinct actor objects in TAG. In those cases, mapping ActorX would be ambiguous because we do not know whether we should map it with actor object ActorY or actor object ActorZ in TAG. This is similar to the “actor ambiguity” described in subsection 4.4.2.

To not add ambiguity to TAG, we did not include those mappings. For our subsequent analyses, we also filtered out the associated IOCs. We encountered this situation 11 times, excluding 6,701 IOCs (0.05%).

#### AUGMENTATION RESULTS


We augmented the TAG mapping with 871 unique actors and 4 countries, nearly doubling the number of actors in TAG. The total number of actors in the augmented TAG is 1,726, and the total number of countries is 41. However, out of the 871 actors added, 833 actors appear in only one vendor feed; they are singletons. Furthermore, 614 of the added actors have a WIP label as a name (e.g., DEV-XXX or TAG-XXX, see subsection 4.4.2). In sum, the majority of the augmented mapping consists of actors with no aliases, and frequently with a WIP label.

## 4.6. RESULTS: ACTOR TRACKING

We now address RQ1: *what is the scope of actor-tracking by vendors?* We operationalize the scope of actor-tracking along several dimensions. First, we analyze the number of actors reported by vendors over time to understand trends in actor-tracking growth and the relative coverage of actors per vendor. We also aim to estimate how much of the overall actor population is reported by vendors. Since no ground truth exists for the complete population, we use TAG as a proxy for the broader actor landscape tracked by entities within the TI community, as reflected by their inclusion in TAG. Next, we examine the overlap in tracked actors between vendors and assess how many additional actors are captured when a new vendor is added. Finally, we focus on actor-specific tracking by pairs of vendors: when two vendors track the same actors, do their IOCs overlap, or does each vendor report different artifacts associated with the actor?

### 4.6.1. TRACKED ACTORS OVER TIME

For each vendor and calendar year, we identify actors first observed in that year (i.e., not seen in earlier years for that vendor). We then compute the cumulative sum of these uniquely observed actors over time. These counts are limited to the periods for which we had access to the IOC feeds (see Figure 4.1). Vendors may have published outside our collection window, so the true number of tracked actors is likely higher.



**Figure 4.2:** Cumulative number of tracked actors per vendor over time, includes WIP actors.


In Figure 4.2, all vendors show steady growth in the number of tracked actors. One vendor, V1, sharply increases the number of tracked actors from 2019 onward. This is primarily due to the use of WIP labels. Our feed for vendor V6 ends in 2022. Table 4.1 summarizes totals per vendor; V1 has the highest number of tracked actors (721), V7 has the lowest (39). However, the totals also reflect the time periods during which the feeds were collected. Outside those periods, additional actors are likely covered.

Across all seven vendors, we identify 1,155 unique actors (565 excluding WIP labels). Table 4.1 depicts the actors per vendor and the number of actors in TAG that are reported on. It shows that not a single vendor reports on more than a fifth of all threat actors listed in TAG. Individual vendors vary greatly in reporting on the number of actors listed in TAG, with V7 reporting on only 2% of all TAG actors, and V1 and V2 reporting on 17%. If we combine the tracked actors of all vendors, then together they report on 293 actors listed in TAG, 34% of all TAG actors (855). In other words, an organization that purchases all seven IOCs feeds, at a cost of USD 1-2 million per year, would still capture only 34% of the threat actor landscape as identified by the industry as a whole.

#### 4.6.2. OVERLAP AND UNION IN TRACKED ACTORS

We determine the overlap and union of tracked actors among vendors. In doing so, we compare the pairwise overlap of tracked actors for two vendors and the union of tracked actors between multiple vendors.

First, we want to know to what extent vendors track the same actors. To do so, we calculate the pairwise overlap of tracked actors among vendors using the Szymkiewicz–Simpson overlap coefficient  $\mathcal{O}$ . We use the overlap coefficient because the sizes of the participating sets (i.e., tracked actors per vendor) are very different, and the overlap coefficient normalizes to the smaller set. Let  $A_v$  denote the set of actors reported by vendor  $v$  under



**Figure 4.3:** Overlap coefficient ( $\mathcal{O}$ ) of tracked actors between vendors.

our augmented TAG mapping (where an actor is associated with a list of names). For each vendor pair  $(i, j)$  we compute  $\mathcal{O}$ , defined as:

$$\mathcal{O}(i, j) = \frac{|A_i \cap A_j|}{\min\{|A_i|, |A_j|\}}.$$

Figure 4.3 shows a heatmap of  $\mathcal{O}(i, j)$  across all vendor pairs (we also include Venn diagrams in Figure C.1 in section C.1.) Across all 21 vendor pairs, the mean overlap coefficient is 0.31, indicating modest overlap. The minimum overlap coefficient is 0.18, and the maximum overlap coefficient is 0.46. Thus, for all pairs, the majority of actors tracked by one vendor are not tracked by the other, and each vendor seems to focus on a mostly different part of the actor landscape.

Second, how much does each additional vendor add in terms of coverage of the actor landscape? In Figure 4.4, we show the number of new actors each new vendor can add to the total set. For each additional vendor, we select the vendor that adds the highest number of additional actors. We show this for actors included in the original TAG mapping and for actors we added to TAG in our augmentation. We see that in both cases, there are diminishing returns. The seven vendors combined cover 293 actors from the original TAG mapping, which is shown in Figure 4.4. This covers 34% of all original TAG actors. The union of augmented TAG actors follows a similar pattern, but does not diminish at the same rate, reaching over 500 unique actors in total.

A 2018 study found that organizations used approximately 6 to 10 different TI feeds

(of which 3-5 are commercial sources) [105]. Thus, most organizations use attributed TI that covers only a minor part of the threat actor landscape. Given the hefty price tags on commercial TI, getting TI that reports on most (TAG) actors would be prohibitively expensive, given that even our seven vendor feeds cover only a third of the TAG actors.

### 4.6.3. IOC OVERLAP FOR JOINTLY TRACKED ACTORS

In subsection 4.6.2, we observed that the overlap in tracked actors is limited. We now turn to the IOC level to assess how comprehensive the visibility of two vendors is when tracking the same actors. Specifically, to what extent do the vendors' IOCs for these actors overlap? As in the previous analysis, we use the overlap coefficient to account for size imbalances between vendors' IOC sets.

For each vendor pair  $(i, j)$  and each actor  $a$  present in both  $A_i$  and  $A_j$ , we form the sets of unique IOCs  $I_{i,a}$  and  $I_{j,a}$ . We compute  $\mathcal{O}(I_{i,a}, I_{j,a})$  by collecting all IOCs per vendor for those actors and computing the intersection of those IOCs between vendors. We use these inputs to compute the overlap coefficient for the pairwise vendor matrix in Figure 4.5.

The average of all pairwise overlap coefficients is 0.10. This shows that across vendors there is little overlap of IOCs. Note that this is calculated only including the IOCs of those actors that the vendors claim to be tracking. Even when vendors track the *same* actor, the IOCs they publish barely overlap. In other words, the visibility of each vendor on an actor is very limited. One potential explanation is the different kinds of telemetry that vendors have, which would be reflected in the indicator types. Figure C.3 in section C.1 shows the number of IOCs per indicator type for each vendor. It is evident that vendors vary in the types of indicators they release. Several vendors provide MD5 hashes most frequently, followed by domains. Only one vendor predominantly provides IPv4 indicators.

These findings align with prior work showing limited shared observables across open [95] and commercial TI sources [25]. The latter study only analyzed the overlap among two vendors that jointly tracked 22 actors. It reported between 2.5% and 4% overlap for two vendors. In our study, based on a more comprehensive dataset, we do find a slightly larger overlap, namely of 10% on average.

Across all pairwise vendor comparisons and indicator types, we observe an average overlap of approximately 10% – i.e., an overlap coefficient of 0.10 – with a minimum of 1% and a maximum of 23%, indicating substantial variability in overlap among vendors. While V1 and V2 have high overlap, they are also the vendors with the longest collection window, depicted in Figure 4.1. However, V3, with a much shorter collection window, also has high overlap with several vendors. This suggests that higher overlap is not merely due to longer collection windows.

## 4.7. RESULTS: ATTRIBUTION AGREEMENT

We now turn to our second research question (RQ2): *How much agreement is there on the attribution of actors and countries across vendors?* We answer this question by looking at IOCs that are observed and attributed by two or more vendors. For example, the domain `evildomain.com` may be reported in two IOC feeds. One vendor may attribute it to Ac-

torA while another vendor to ActorB. For simplicity, we will refer to these as ‘co-observed IOCs’. However, note that ‘co-observed’ also means that both vendors attributed the IOC, possibly to different actors. At the same time, most IOCs in the feeds are not attributed by anyone or only by a single vendor. Those are excluded from this analysis. As described in subsection 4.4.3, we measure the agreement in attribution using the Krippendorff’s  $\alpha$  metric for IRR. For each co-observed IOC, we only consider the vendors that ‘voted’ (i.e., attributed) the actor. So there are no ‘missing votes’. Using this approach, though, we cannot compare specific vendors, only overall cross-vendor agreement.

As discussed in section 4.4, we use observation periods (OP) to ensure that the different observations of the same IOC happen within a certain timeframe. For example, for OP=7, we analyze IOCs that were observed by two or more vendors within a period of 7 days. If the observations are very far apart, they are less likely to belong to the same attacker and, thus, less likely to be both correctly attributable to that actor. So this would bias the agreement metrics against the vendors. We want to be conversative, so we only compare IOCs that are closer in time. To understand if the observation period impact the agreement, we analyze our data for different periods: 7, 14, 30, and 60 days. In subsection 4.7.1, we describe the number of co-observed IOC overall and per observation period. First, we examine agreement over indicators for co-observed IOCs that include WIP label attributions (described in subsection 4.4.2). However, we do not consider WIP attributions to be real *actor* attribution. Additionally, they lead to disagreement because the label is vendor-specific. Therefore, the subsequent analyses consider attributions excluding WIP labels. Next, we investigate the set of co-observed IOCs by two and three vendors. The number of observations by four vendors (see Table C.6) was so low that we cannot derive any conclusions. No IOCs were co-observed by five or more vendors. Furthermore, in subsection 4.7.2, we try to understand disagreement for MD5 hashes by looking at the types of software, obtained from VirusTotal (VT) [263]. In subsection 4.7.3, we determine how specific vendors relate to each other when co-observing IOCs by computing the pairwise agreement between specific vendors for different observation periods. This analysis does not include IOCs with WIP labels.

## 4

#### 4.7.1. CO-OBSERVED IOCS

IOCs that are co-observed and attributed by two or more vendors are rare. If we do not apply any limitation on how far apart the observations took place, then 135,987 IOCs (1% of all attributed IOCs, including WIP labels) are co-observed. Excluding WIP labels, there are 112,625 IOCs. The distribution of co-observed indicator types per vendor is depicted in Figure 4.6. It shows that hashes and domains are most frequently co-observed, URLs the least. Also, vendors V1 and V2 account for a large portion of co-observed IOCs.

As soon as we select a observation period, the maximum distance in time between the observations of different vendors is set to 7, 14, 30, or 60 days. The smaller the OP, the lower the number of co-observed IOCs. Table 4.3 shows the number of co-observed IOCs for each observation period (OP) for actor and country attribution. As expected, the most common case is two vendors co-observing an IOC within the 60-day OP, namely 12,326 IOCs (15,822 exc. WIP labels).

For the country attribution, there are 11,428 IOCs (12,519 exc. WIP labels) for two vendors over a 60-day OP. The maximum number of vendors that co-observe an IOC is

four, depicted in Table C.6. No IOCs are co-observed by five or more vendors.

#### CO-OBSERVATIONS OF IOCs INCLUDING WIP LABELS

Table 4.3 reports the overall agreement for IOCs *including* WIP labels (type=ALL (INC . WIP)). We find poor agreement across observation periods:  $\alpha$  ranges from 0.56-0.61. Agreement among vendors is the highest when the observation period is shorter ( $OP \leq 14$ ), and declines a bit ( $\alpha = .56$  and  $\alpha = .57$ ) when the observation period is increased to 30 and 60 days. This decrease, however, is small, so the chosen observation period does not seem to have a substantial impact on the level of agreement. By contrast, country attribution has high agreement across all observation periods and indicator types ( $\alpha$  ranges from 0.89-0.92). These observations are similar for co-observed IOCs by 3 vendors. Vendors are, thus, likely to disagree over actors but not over the origins of the attacks.

As stated before, we do not consider WIP label IOCs strong attributions, and neither do vendors [259]. We focus all analysis below on agreement for attributions *excluding* WIP labels.

#### CO-OBSERVATION BY TWO VENDORS

Table 4.3 shows that IOCs co-observed by two vendors are most prevalent. For type ALL, we find moderate agreement across observation periods:  $\alpha$  ranges from 0.76-0.81. Agreement among vendors is highest when the observation period is shortest ( $OP \leq 7$ ), and marginally declines ( $\alpha = .78$  and  $\alpha = .76$ ) when the observation period is increased to 30 and 60 days. This is not surprising, as attacker infrastructure more likely to have been abandoned and changed hands over time [22]. The chosen observation period does not seem to have a substantial impact on the level of agreement. Conversely, country attribution approaches near-perfect agreement across observation periods and indicator types ( $\alpha$  ranges from 0.92-0.95).

Agreement among vendors is lowest for IPv4 IOCs ( $\alpha$  ranges from 0.65-0.71). Also, IPv4 indicators are co-observed less often. The agreement is moderate ( $\alpha$  ranges from 0.65-0.71) and remains somewhat constant across observation periods. This relative disagreement is perhaps not surprising as IPv4 addresses are believed to be transient attacker infrastructure with high churn [22, 53].

URL indicators are the least observed IOCs. Similar to MD5 indicators, they have a satisfactory and stable level of agreement for all observation periods ( $\alpha > 0.80$ ). Contrary to IPv4 addresses, URLs and hashes tend to be of a more permanent nature, which seems reflected in the level of agreement. For this reason, operational security (OPSEC) best practices dictate that files and payloads should not be reused [27]

DOMAIN indicators have the least stable agreement across observation periods. These IOCs have high agreement at the shortest observation period, which consistently declines ( $\alpha$  ranges from 0.84-0.76). This may be due to infrastructure churn when domains are identified as malicious.

Overall, we find that the overall level of agreement is moderate and doesn't seem to be substantially impacted by a longer observation period. This finding somewhat undermines the idea that organizations can use attributed IOCs to detect and respond to particular adversaries effectively, particularly for IPv4 indicators. While agreement on the actor is not poor, only agreement on the country consistently scores high.

## THREE-VENDOR CO-OBSERVATION

A third vendor also observing an IOC within the same observation period is rare. In total, there are 730 cases at most (see row (ALL), OP=60, column 3 vendors). The number of co-observed IOCs is particularly low for observation periods of 7 and 14 days across the various indicator types. For country attribution, the number of IOCs is equally low, but the  $\alpha$  values remain high compared to the two-vendor agreement.

When analyzing the alpha by indicator type, we find that the DOMAIN indicator results in poor agreement, despite several observations. By contrast, the MD5 indicator has very high agreement. This may be due to the permanent nature of hashes. The IPv4 indicator has better agreement for observation periods of 14 and 30 days ( $\alpha$  ranges from 0.78-0.81), but the number of observations is low. Once the observations increase (OP=60), the agreement is poor. This is similar for the URL indicator, where the longest observation period yields poor agreement.

The low number of observations seems to result in less stable alphas. While the agreement is higher, the confidence intervals are wider. So the net effect for observations by three vendors is not higher agreement. For actor attribution, this undermines the idea that having multiple TI sources would make attribution easier for customers. By contrast, country attribution consistently results in high levels of agreement, regardless of the number of vendors.

## ATTRIBUTION ANALYSIS OF SECONDARY IOCS

Secondary IOCs include types such as Yara and Snort rules, filenames, MIME types, emails, X.509 certificate details, auth hashes, crypto addresses, passwords, user agents, and imphashes (55 in total). These are largely vendor-specific. We observe only 3,706 secondary IOCs co-observed by two vendors (3% of all co-observed *primary* indicators), with none observed by more than two vendors. All were exclusively shared between V3 and V4, likely reflecting their vendor-specific nature. Table C.7 in section C.1 summarizes agreement levels.

We find lower sample sizes ( $n = 93-375$ ), low actor attribution agreement ( $\alpha = 0.29-0.40$ ), and wider confidence intervals. Country attribution also has low  $n$  but perfect agreement ( $\alpha = 1$ ). These results suggest that vendor-specific IOCs lead to lower actor attribution agreement. As discussed in ??, we hypothesize that such IOCs stem from vendor-specific telemetry (e.g., customer base and network vantage points), potentially introducing attribution biases.

## 4.7.2. DISAGREEMENT ON MD5 INDICATORS

What are the sources of disagreement between the vendors? This cannot be answered without inside knowledge of not just one vendor, but all of them. This is not available to anyone, not even the vendors themselves. Moreover, in the absence of ground truth, the disagreement also cannot be traced by to ineffective practices of one vendor over another.

So we can only speculate about the sources of disagreement. For most indicator types, it is not difficult to understand why vendors might reach different conclusions. Indicators like DOMAIN, URL, IPv4 have transient properties. IP addresses are re-assigned regularly, domains might be compromised assets that change hands again later. It is

more surprising that also file hashes lead to different attributions. The hashes point unambiguously at very specific artefacts that are frozen in time. Why would vendors not have higher agreement rates there than for the more transient indicators?

Using VirusTotal (VT), we took a closer look at some of the files associated with the MD5 hashes. are permanent and unambiguous. We looked up 867 co-attributed hashes for a 7-day observation period for two vendors. For this set, 746 hashes resulted in agreement, 121 in disagreement. For the 746 agreement hashes, VT returned 532 results. For the 121 disagreement hashes, VT returned 68 results. In total, VT had results for 69% of the hashes.

In Table 4.4, we depict the counts for agreement and disagreement for several file types. In our view, there are two possible explanations for these discrepancies. First, some vendors may mislabel certain files. This is problematic because if a company relies on the corresponding feed, it could lead to false positives. Second, some files may have a dual-use nature: they can be used for benign purposes under certain conditions, but for malicious purposes under others. For example, various studies looked at how attackers are ‘living off the land’ and appropriating legitimate resources for malicious activities [19, 188].

Other files types, like Microsoft Excel spreadsheets or Office documents, are more likely associated with initial access attempts. They might contain exploit code that can be correlated with artefacts from other attacks. Yet they still reveal less about the attackers than the full malware packages they rely on later in a multi-stage attack. So it provides less evidence on which to base attributions. Even though the object itself is stable and unambiguous, the attribution process still operates under significant uncertainty.

### 4.7.3. PAIRWISE AGREEMENT AMONG VENDORS

The previous subsections examined overall attribution differences among vendors, i.e., they were not vendor-specific. We want to know if overall disagreement perhaps arises from specific vendors that mostly disagree with others. Therefore, we analyze agreement differences among specific vendors by computing the pairwise agreement, described in subsection 4.4.3. For the pairwise agreement, we compare matches (i.e., the same attributions) against all co-observed IOCs, excluding IOCs with WIP labels.

Figure 4.7 shows pairwise agreement among vendors for co-observed IOCs across multiple observation windows. We find that some vendors agree more than others across observation periods, but generally speaking, disagreement is distributed across the vendor population. In some periods, V2, V3, and V6 form a consistently high-agreement cluster, whereas V1 and V5 show comparatively lower agreement with others. V7 frequently behaves as an outlier, diverging from the remaining vendors; a potential explanation is its regional focus. Agreement generally weakens as the observation window widens, and by 60 days, no vendor pair exhibits perfect agreement.

## 4.8. DISCUSSION

We find that vendor coverage of tracked threat actors is limited, with little overlap in attributed IOCs across vendors, even when they track the same actors. Prior work has similarly highlighted challenges in coverage, overlap, timeliness, and accuracy [25, 126,

136, 152, 243].

The low overlap suggests that visibility into attacker activity is largely shaped by vendor-specific telemetry, influenced by factors such as customer base and network vantage points. As a result, each vendor captures only a partial view of the attack landscape. Whether combining these partial views yields a near-complete picture for defenders, though, remains an open question.

We now reflect on the implications of our findings for the three use cases of attributed TI, as outlined in the Introduction.

#### 4.8.1. IMPLICATIONS

**Threat-led defense strategies.** For defenders, the customers of high-end TI, our findings help reduce the information asymmetry they face vis-à-vis vendors. It is very difficult for a customer to evaluate the properties and quality of TI. For the issue of attribution, we now have at least an empirical assessment of what the market delivers.

The fragmented and inconsistent nature of commercial threat attribution poses a direct challenge to threat-led defense strategies. Our findings show that each vendor captures only a fragment of the actor landscape as represented by MISP TAG. This sparsity means organizations relying on one source will have blind spots in their defenses, potentially missing adversaries that other feeds identify. Even combining all seven top-tier feeds only covered roughly one-third of known actors, indicating that defenders cannot assume coverage of the full threat spectrum. In practice, this undermines the common guidance that enterprises should ‘know the adversaries targeting them’ and might lead to misallocated investments or gaps in monitoring.

Moving to the IOCs, we found that most of them are not attributed to any actor. Where there is attribution, it is mostly by a single vendor. Even when vendors claim to track the same actors, they find mostly different IOCs. When we look at the set of attributed IOCs that are reported by more than one vendor within the same observation period, we find only modest agreement in the attribution.

Still, this is the most generous comparison we can make for vendors, and some level of disagreement still remains. This also implies that the singleton attributions, which make up the bulk of the attributed data, have a decent probability of being wrong.

**Legal Exposure and Sanctions Screening.** Attribution inconsistency can have serious legal and compliance ramifications. Many organizations today use threat intelligence to check whether a breach or a ransomware attack involves actors under legal sanctions or other restrictions. For example, the U.S. Office of Foreign Assets Control (OFAC) maintains sanctions against certain cyber adversaries, and companies must avoid transactions (like ransom payments) that benefit those groups. If threat feeds provide sparse or disjoint coverage, a victim organization might fail to recognize that an attacker is a sanctioned entity simply because their intel vendor labels the group differently or not at all.

The high agreement on country-level attribution (Krippendorff’s  $\alpha > 0.9$ ) offers some consolation – at least geographically, threat intelligence largely concurs. This consistency means that if an organization knows, for instance, that any state-sponsored attacker from Country X must trigger certain legal protocols (government notifications, export control checks, etc.), the country tag from a feed is likely reliable. However, most

legal obligations (like sanctions) are tied to specific named groups or individuals, not just countries.

**Geopolitical Attribution and State Response.** In the context of national security and state-level responses, our findings have dual implications: they highlight a challenge at the actor granularity but a reassuring alignment at the country level. Governments often rely on both private intelligence and their own agencies to attribute cyber incidents to perpetrator groups or sponsor states. If an incident is labelled as the work of Group A by one prominent vendor but Group B by another, international partners might each cite different intelligence sources, muddying the waters of attribution. Adversary nations can exploit these discrepancies for propaganda or denial: a state accused of an attack could point to alternative attributions from respected vendors to cast doubt on the accusation. Thus, inconsistent actor attribution risks undermining the credibility of geopolitical claims. On the other hand, our data shows that when it comes to attributing attacks to a sponsoring country, strategic responses (diplomatic *démarches*, sanctions, or even retaliatory cyber operations) can be decided with greater confidence. In practice, governments seldom rely on vendor reports alone for such decisions, but these reports shape public narratives and ally coordination.

#### 4.8.2. RECOMMENDATIONS

We outline recommendations for different stakeholders in the threat intelligence ecosystem:

**TI Consumers (Defenders).** Organizations might be better off adopting a broader, technique-focused approach to defense rather than focusing on actors. TTPs might not be as helpful in attribution as is often assumed [254], but they do provide insight into what techniques might be deployed against their defenses – for example, using MITRE ATT&CK patterns common to multiple actors. If resources allow, use multiple intelligence sources and cross-verify actor attributions rather than relying on one feed. Finally, ensure that compliance teams cross-check any identified actor against official sanction or watch lists using all known aliases, not just the names of their own vendors.

**TI Vendors.** There are benefits to collaborating on common actor identifiers or at least transparently mapping actor names to industry-recognized aliases [154].

While there are good reasons for initially using vendor-specific aliases [213, 226], the proliferation of over 10 synonyms for the same threat group is a disservice to defenders. Eventually, aliases should be merged.

Vendors should also clarify their focus areas and tracked actors, so clients understand which parts of the actor landscape might not be covered by a given feed.

Improving coverage likely depends on information sharing in the industry, similar to how the anti-virus vendors pooled their malware samples. To some extent, this is undoubtedly already happening, as reflected in vendor reports that explicitly link to the attributions of other vendors. It suggests intelligence sharing via trusted channels [216]. Yet the ability to more comprehensively cover the threat landscape still seems out of reach. To stick with the parallel of the AV industry: any AV solution worth its salt needs to defend against a much wider range of malware than the samples that the vendor is able to collect. So all products improve from pooling samples. There are free-rider risks in such an arrangement [49, 164], but overall it has become the norm. The benefits of

sharing have turned out to be stronger than the incentive to monetize exclusive, proprietary data. The TI industry should consider going the same route.

**Policymakers and Regulators.** Government and international bodies can work with the security industry to maintain an up-to-date public taxonomy of threat groups (building on efforts like MISP TAG or MITRE's group listings) that includes known aliases and supporting evidence of equivalence. This would give organizations a reference point to reconcile vendor information, especially for legal and diplomatic purposes. Finally, governments include intelligence agencies. These have unique methods to support actor attribution. Governments can support collaboration in this nexus. While intelligence agencies cannot cover the width of the private sector, sharing with the private sector would increase the protection of civil entities relying on those commercial products.

## 4

### 4.8.3. LIMITATIONS

Our study has several limitations. First, we do not have ground truth data on actors. This prevents us – and in most cases, everyone else – from knowing whether an attribution is correct. To cope with this limitation, we have focused on (dis)agreement among vendors as a proxy for accuracy.

Second, as we rely on our partner organization, we work with the set of TI vendors that they have contracted with. This makes our findings biased towards a single enterprise partner. Moreover, while these are top-tier TI vendors, as indicated in the fees for their products, no set of seven vendors could claim to represent the bulk of the market. Many firms produce TI, either as their main focus or as a by-product of other security services. Specific numbers across buyers and TI feeds may vary. However, given the absence of scientific reporting on attribution and the unavailability of expensive commercial TI feeds to researchers, we believe that, despite potential biases, our study still yields valuable insights.

Third, the country agreement we reported may be inflated, as we derived country labels via the TAG mapping, without checking all industry reports. However, during the process of extending TAG, we checked many original TI reports and found no disagreement on the country.

Fourth, we decided to exclude IOCs that led to ambiguous mappings. The actual disagreement of vendors is therefore likely to be higher. We acknowledge this, but we believe that attribution, by its very nature, is mired in uncertainty. We wanted to be conservative and err on the side of agreement by focusing the comparison on the most favorable conditions for reaching agreement.

Fifth, vendor V2 provides many IOCs, many more compared to the other vendors. This introduces a bias where we measure the agreements of other vendors with vendor V2. Nonetheless, other vendors also contribute greatly to the set of co-observed IOCs, so we consider the current approach valuable.

Finally, our augmentation process relies on vendor reports and is a manual task. This may propagate vendor-specific biases and annotator biases. As described in section 4.3, vendor reports play a part in the attribution process. We treat this as a validation of our approach and believe the bias to be limited.

#### 4.8.4. FUTURE RESEARCH

This study scratches the surface of actor attribution in TI. We provide three directions for future research. First, we investigated agreement among *commercial* TI vendors. Future work could focus on comparing attribution among public sources and commercial sources versus public sources. One challenge in this research direction is the mapping (process) for actors. Public sources may not have the authoritative reporting that commercial TI vendors provide. Second, we made an initial attempt at explaining root causes for disagreement by looking at the software types for hashes. New work might examine what other root causes underlie disagreement. Mixed-methods studies that collect and analyze quantitative data and interviews may provide valuable insights. Lastly, future work could combine the IOCs from multiple organizations to obtain a broader understanding.

### 4.9. CONCLUSIONS

We have provided an empirical view on the attribution of 13.5M IOCs by seven leading TI vendors. The gaps we uncovered have impact on the use cases for attributed TI. We explored those implications and reflected on recommendations to improve the status quo.

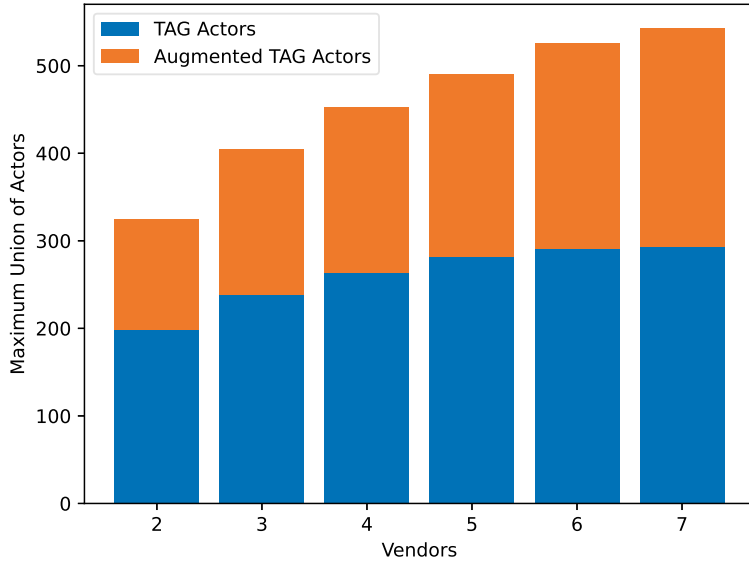


Figure 4.4: Maximizing actor union between vendors overall and in TAG, excluding WIP actors.





Figure 4.5: Overlap coefficient (C) for IOCs associated with actors that are tracked by both vendors.

**Table 4.3:** Actor and country attribution agreement by indicator type and observation periods (OP). Values show Krippendorff's  $\alpha$  with 95% CI and sample size ( $n$ ). All types exclude WIP label IOCs, except ALL (inc. WIP).

Type	OP (Days)	2 vendors		3 vendors	
		Actor	Country	Actor	Country
ALL (inc. WIP)	7	0.60 [0.59,0.61] ( $n=3,530$ )	0.89 [0.88,0.91] ( $n=2,761$ )	0.50 [0.42,0.60] ( $n=47$ )	0.64 [0.58,0.72] ( $n=44$ )
ALL (inc. WIP)	14	0.61 [0.60,0.62] ( $n=5,599$ )	0.92 [0.91,0.93] ( $n=4,612$ )	0.68 [0.63,0.73] ( $n=144$ )	0.89 [0.85,0.95] ( $n=129$ )
ALL (inc. WIP)	30	0.56 [0.55,0.57] ( $n=10,089$ )	0.91 [0.90,0.92] ( $n=7,729$ )	0.68 [0.64,0.71] ( $n=344$ )	0.81 [0.76,0.86] ( $n=322$ )
ALL (inc. WIP)	60	0.57 [0.56,0.57] ( $n=15,822$ )	0.89 [0.89,0.90] ( $n=12,519$ )	0.67 [0.65,0.69] ( $n=851$ )	0.85 [0.82,0.88] ( $n=777$ )
ALL	7	0.81 [0.80,0.82] ( $n=2,748$ )	0.94 [0.93,0.95] ( $n=2,487$ )	0.71 [0.49,1.00] ( $n=35$ )	1.00 [1.00,1.00] ( $n=35$ )
ALL	14	0.79 [0.78,0.80] ( $n=4,524$ )	0.95 [0.94,0.96] ( $n=4,201$ )	0.83 [0.76,0.92] ( $n=120$ )	1.00 [1.00,1.00] ( $n=120$ )
ALL	30	0.78 [0.77,0.79] ( $n=7,642$ )	0.95 [0.94,0.95] ( $n=7,060$ )	0.81 [0.76,0.86] ( $n=295$ )	0.97 [0.95,1.00] ( $n=295$ )
ALL	60	0.76 [0.76,0.77] ( $n=12,326$ )	0.92 [1.00,1.00] ( $n=11,428$ )	0.76 [0.74,0.79] ( $n=730$ )	0.95 [1.00,1.00] ( $n=729$ )
MD5	7	0.83 [0.81,0.85] ( $n=867$ )	0.97 [0.96,0.99] ( $n=695$ )	1.00 [1.00,1.00] ( $n=11$ )	1.00 [1.00,1.00] ( $n=11$ )
MD5	14	0.83 [0.82,0.85] ( $n=1,449$ )	0.98 [0.97,0.99] ( $n=1,234$ )	1.00 [1.00,1.00] ( $n=55$ )	1.00 [1.00,1.00] ( $n=55$ )
MD5	30	0.84 [0.82,0.85] ( $n=2,486$ )	0.98 [0.97,0.98] ( $n=2,158$ )	0.99 [0.98,1.00] ( $n=132$ )	1.00 [1.00,1.00] ( $n=132$ )
MD5	60	0.83 [0.82,0.84] ( $n=3,753$ )	0.96 [0.95,0.96] ( $n=3,282$ )	0.90 [0.85,0.95] ( $n=314$ )	0.90 [0.85,0.95] ( $n=314$ )
URL	7	0.87 [0.83,0.90] ( $n=275$ )	0.99 [0.98,1.00] ( $n=243$ )	1.00 [1.00,1.00] ( $n=1$ )	1.00 [1.00,1.00] ( $n=1$ )
URL	14	0.89 [0.86,0.91] ( $n=407$ )	0.99 [0.98,1.00] ( $n=365$ )	1.00 [1.00,1.00] ( $n=1$ )	1.00 [1.00,1.00] ( $n=1$ )
URL	30	0.82 [0.80,0.85] ( $n=625$ )	0.99 [0.98,1.00] ( $n=556$ )	1.00 [1.00,1.00] ( $n=9$ )	1.00 [1.00,1.00] ( $n=9$ )
URL	60	0.87 [0.85,0.89] ( $n=951$ )	0.99 [0.98,1.00] ( $n=844$ )	0.70 [0.61,0.79] ( $n=44$ )	1.00 [1.00,1.00] ( $n=44$ )
DOMAIN	7	0.84 [0.83,0.87] ( $n=1,236$ )	0.98 [0.97,0.99] ( $n=1,189$ )	0.48 [0.45,0.49] ( $n=22$ )	1.00 [1.00,1.00] ( $n=22$ )
DOMAIN	14	0.78 [0.77,0.80] ( $n=1,916$ )	0.99 [0.98,1.00] ( $n=1,864$ )	0.66 [0.55,0.78] ( $n=59$ )	1.00 [1.00,1.00] ( $n=59$ )
DOMAIN	30	0.79 [0.78,0.80] ( $n=3,224$ )	0.98 [0.97,0.99] ( $n=3,103$ )	0.63 [0.56,0.70] ( $n=138$ )	0.94 [0.86,1.00] ( $n=138$ )
DOMAIN	60	0.76 [0.75,0.77] ( $n=5,554$ )	0.94 [0.93,0.95] ( $n=5,351$ )	0.65 [0.61,0.70] ( $n=294$ )	0.94 [0.88,1.00] ( $n=293$ )
IPV4	7	0.65 [0.61,0.69] ( $n=278$ )	0.93 [0.90,0.96] ( $n=268$ )	0.00 [0.00,0.00] ( $n=1$ )	1.00 [1.00,1.00] ( $n=1$ )
IPV4	14	0.71 [0.68,0.74] ( $n=601$ )	0.96 [0.95,0.98] ( $n=587$ )	0.78 [0.55,1.00] ( $n=5$ )	1.00 [1.00,1.00] ( $n=5$ )
IPV4	30	0.66 [0.63,0.68] ( $n=1,070$ )	0.96 [0.95,0.97] ( $n=1,008$ )	0.81 [0.68,1.00] ( $n=16$ )	0.93 [0.83,1.00] ( $n=16$ )
IPV4	60	0.66 [0.64,0.68] ( $n=1,698$ )	0.94 [0.93,0.95] ( $n=1,584$ )	0.65 [0.59,0.70] ( $n=78$ )	0.96 [0.93,1.00] ( $n=78$ )



**Figure 4.6:** Distribution of actor attributed IOCs co-observed by 2+ vendors, excluding WIP labels.

**Table 4.4:** Filetypes for which vendors both agreed and disagreed in attribution

File type	Agree	Disagree	Ratio (%)
MS Excel Spreadsheet	6	11	183
Outlook	1	1	100
Win32 DLL	38	22	58
ISO image	3	1	33
Windows shortcut	61	10	16
Win32 EXE	84	15	18
VBA	9	2	22
Powershell	9	1	11
HTML	109	3	3
Office Open XML Document	44	1	2

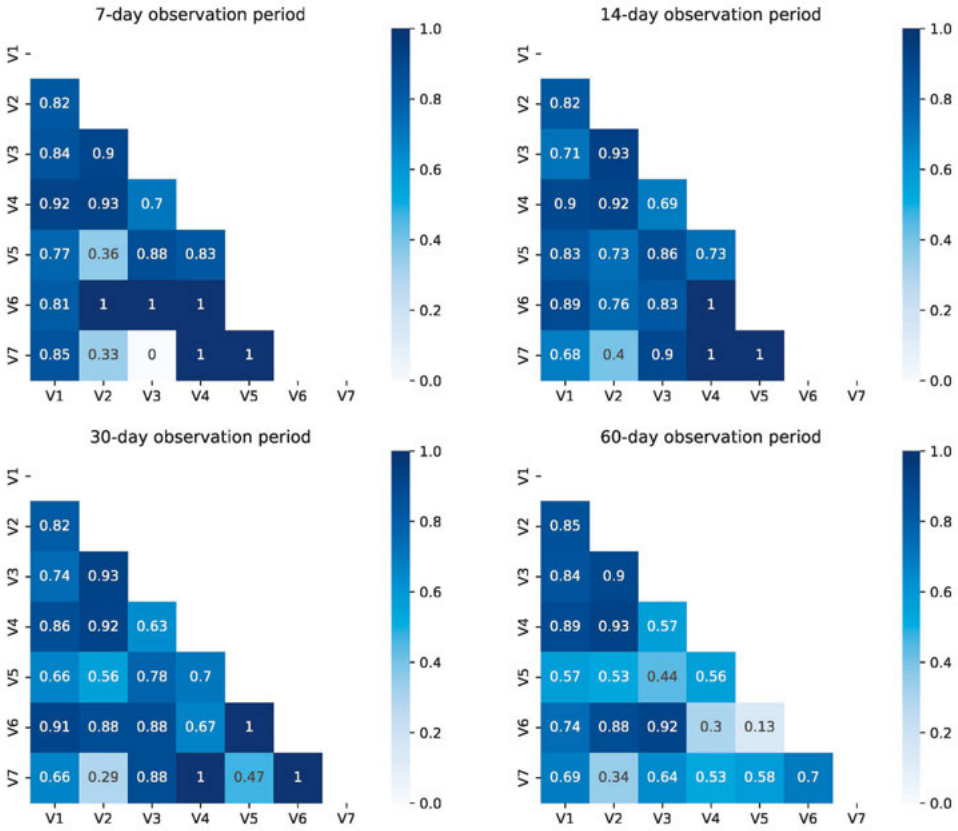


Figure 4.7: Pairwise vendor agreement on actor attribution for co-observed IOCs.



# 5

## CONCLUSION

This dissertation studied security measures of municipalities to address cyber threats. A total of three peer-reviewed studies have been presented in Chapters 2 to 4 and were all aimed at answering the following research question:

*How can municipalities improve security measures to address cyber threats?*

In this final chapter, we summarize the empirical findings from Chapters 2-4 to answer our main research question. Then, we proceed to reflect on the empirical findings. Next, we consider the impact of our work and its implications for municipal cybersecurity governance. Finally, we propose future research directions that build on the findings in this dissertation.

### 5.1. EMPIRICAL FINDINGS

Chapters 2 to 4 presented the results from the studies into the security measures of municipalities. The main conclusions of the studies are summarized in the following paragraphs and centre around the research gaps identified in section 1.2.

#### 5.1.1. CHAPTER 2 - PATCH BEHAVIOR AND MANAGEMENT OF VULNERABLE HOSTS

In Chapter 2, we asked what explains patching behavior and management for vulnerable hosts. Using similar detection techniques as O'Hare et al. to determine vulnerable hosts [186], we scanned approximately 300,000 IPv4 hosts of 322 municipalities and observed 3,402 responding hosts, 578 hosts with a version service. From these 578 hosts, 101 unique CPEs were derived, of which 70 contained 1 or more CVEs. We observed 70 vulnerable services at 154 unique hosts in 94 different municipalities. During the interviews of 29 practitioners, we validated 15 vulnerable hosts running 4 vulnerable services: Nginx, OpenSSH, Apache, PowerDNS. We also validated 27 CPEs that ran Microsoft services; however, they did not contain CVE information. In total, we found that the 42 detected CPEs were indeed correct. Furthermore, we observed 33 vulnerable hosts in total,

6 hosts were only observed by the respective municipality, 18 hosts were only observed by the CERT, and 9 hosts were observed by both the municipality and the CERT. This discrepancy illustrates a responsibility gap around shadow IT and other systems outside the reach of the IT organizations. Many vulnerable hosts fell into this responsibility gap. For 18 hosts, the system administrators were not aware of the vulnerable systems, and they did not consider these systems their responsibility. For 6 hosts that the administrators considered their responsibility, the CERT was unaware of them. Municipal systems registered at the CERT often turned out to be incorrect and/or out-of-date. For the 15 hosts that sysadmins acknowledged as being responsible for, we found 3 explanations for not patching the system: admins were not aware of the system (9 hosts), admins were unable to patch the system (3), and the systems were being retired (3). We also found a fourth explanation for non-patching behavior that we could not measure externally but did surface as a potential option for sysadmins during the interviews, a vulnerable system could not be patched but could be shut down, i.e., disconnected, to prevent harm.

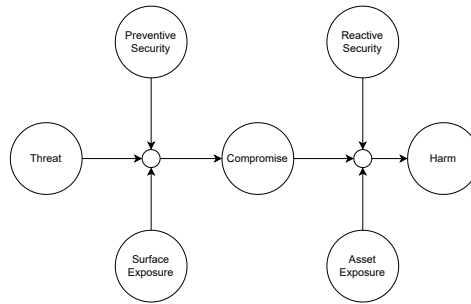
## 5

### 5.1.2. CHAPTER 3 - FUNCTIONING OF INSTITUTIONAL SUPPORT STRUCTURES FOR MUNICIPAL INCIDENT PREVENTION AND MITIGATION

This chapter aimed at evaluating the functioning of institutional support structures for incident prevention and mitigation at municipalities, specifically sector Computer Incident Security and Response Teams (CSIRTs). We engaged with 26 practitioners from three stakeholder groups—i.e., governing bodies, sector CSIRT staff, and constituents—to learn their perspectives on the expectations of the Informatiebeveiligingsdienst (IBD), the sector CSIRT for municipalities. Additionally, we identified the challenges in providing those services. We found that practitioners' daily practices are shaped by three dynamics: resources, legitimacy, and dependency. We found that the actual provided services often do not align with the expectations of constituents, especially around incident response. Furthermore, we observed several strategic challenges for sector CSIRT practitioners. These include: a diverse constituent population and maturity, trust issues, and service-specific organizational dependencies. Finally, we evaluated the functioning of the vulnerability notification service, a key service of the sector CSIRT. We conducted a historical analysis of data of the vulnerability notifications that were sent to the constituents between 2015-2024. Our analysis uncovered a systematic problem where many of the notifications that should have been sent to the constituents never reached the IBD-CSIRT, because the national CSIRT did not forward them. Moreover, the practitioners did not detect this problem, signalling a missing feedback mechanism.

### 5.1.3. CHAPTER 4 - QUALITY EVALUATION OF ATTRIBUTION IN COMMERCIAL TI

In the third study in Chapter 4, we sought to systematically investigate the attribution of IOCs in high-end commercial threat intelligence feeds, providing a longitudinal comparative analysis across 13.5 million IOCs of seven vendors. We validated and augmented a leading industry actor mapping, MISP Threat Actor Galaxy (MISP TAG or TAG), which contains 855 actors. Our validation showed TAG has high accuracy when checked against a deconflicted ground-truth mapping from Crowdstrike and Microsoft, and against vendor TI reports for a random sample of 50 actor names. We augmented the TAG mapping



**Figure 5.1:** The social-technical causal model as a structured way to understand municipal security challenges.

to a total of 1,726 actors, using the vendor reports, adding 871 new actors. Next, we found that vendor coverage of threat actors is limited. Each vendor covers only a fraction of the actors in TAG, between 2-17%. The union of all seven feeds, which would cost USD 1-2 million per year, can track just 34% of the actors from TAG. Next, we found that just 1% of all attributed IOCs are observed by more than one vendor. Even when vendors state that they track the same actors, the IOCs they report for those actors barely overlap. When vendors do observe the same IOC, their attribution often disagrees: Krippendorff's  $\alpha$  ranges from 0.57-0.62 for two vendors, even when the observations are just a week apart in time. Last, agreement is higher when we remove 'work-in-progress' attributions:  $\alpha$  ranges from 0.78-0.82, across observation periods for two vendors. Agreement on country attribution is consistently high, ranging from 0.92-.95 for two vendors. We identify implications for security practices around the use cases of attributed TI.

5

## 5.2. REFLECTIONS ON FINDINGS

The empirical findings from these studies present a new perspective on the realities of municipal security measures. In this section, we reflect on how these findings may ultimately help municipalities strengthen their defenses against cyber threats.

First, the studies show that municipal cybersecurity is, in large part, an organizational problem rather than a purely technical challenge. Chapters 2 and 3 showed that the effectiveness of the vulnerability notification mechanism is undermined because of unclear asset responsibilities, poor asset management, out-of-date asset registrations, an opaque methodology for processing vulnerability data, and limited resources for security within organizations. Instead, the findings show that municipalities can only improve their security posture when governance, ownership, and accountability are clearly aligned with technical processes. Additionally, the findings highlight the need and importance of interdisciplinary research, as they reveal shortcomings in security measures that would not surface with solely technical or purely social science research. The studies reposition municipal cybersecurity as a problem of coordination and visibility rather than of technology adoption. This insight aligns with and extends prior work on the socio-technical nature of cybersecurity governance [10, 125, 235, 274], showing empirically how institutional design ultimately shapes technical effectiveness.

Second, asset management is a precondition for several security measures. For ex-

ample, in the Netherlands, the BIO regulatory framework for governments, identifies three “basisbeveiligingsniveaus” (BBNs), i.e., basic security levels for government systems. For each security level, certain controls are required or recommended. However, our findings suggest that municipal asset management is problematic and undermines the use of any security framework. Without knowing its assets, a municipality cannot adequately set security levels and controls for its systems.

On a more abstract level, Chapters 2 and 3 reveal that asset management is more than a tedious job of IT staff. Instead, asset management emerges as a governance function rather than a maintenance task; without visibility on systems, neither notification chains nor CSIRT services can operate effectively. These mechanisms fail if assets are not known or incorrectly attributed. Indeed, in Chapter 2, the findings show that municipal assets registered with the CSIRT were out-of-date. Furthermore, the study showed that knowing what system is vulnerable is not enough; it is also vital to know who is responsible for it. Moreover, during the interviews, practitioners indicated that asset management remains a problem for them: assets are fleeting and may go unreported to the IT team, leading to shadow IT. This supports earlier work showing that unregistered IT infrastructure is a widespread challenge [205]. For municipalities, the problem is magnified by the diversity of services they deliver, which creates opportunities for shadow IT to emerge across the organization. As IT adoption continues to grow within municipalities, they become increasingly dependent on those systems that run many different organizational processes. In turn, the attack surface, and thus risk, of the municipality grows. This development maps to an increase in the Surface Exposure factor in the causal model in Figure 1.1. Moreover, as many services move to Software-as-a-Service (SaaS) solutions, responsibilities become even more diffuse, and municipalities lose the ability to act on the infrastructure that provides their services, highlighting a strong interdependence on commercial service providers. Given these developments, asset management problems are likely to worsen, further undermining essential security measures such as patching.

While the root causes for asset management challenges are organizational, some potential solutions may have a technical component to make organizational boundaries and system responsibilities more explicit. First, many modern External Attack Surface Management (EASM) solutions provide automated asset discovery. In the UK, the NCSC ran trials with EASM systems at constituent organizations and their findings reveal that the automated asset discovery was often “good enough” [169]. The trials also showed that those EASM systems identified legacy domains that practitioners thought had been decommissioned. EASM systems also identified new services that the cybersecurity teams didn’t know about. Those findings map closely to the factors for non-patching behavior in Chapter 2. Such tools illustrate that technical innovation can mitigate, though not replace, the organizational clarity on which vulnerability management depends.

Similarly, asset management may also be improved with modern registration tooling. The Dutch NCSC developed a central asset registration platform for organizations across sectors. Instead of constituents registering assets to the sector who, in turn, report changes to the NCSC, there is now a single online location for asset registration. This service sets a norm for organizations across sectors: register and update your assets. While

this does not solve organizational challenges, it will provide a step forward in improving municipal security posture.

Furthermore, the study in Chapter 2 finds that patching remains a challenging process for municipalities. The first difficulty lies in determining who is responsible for a system. Once responsibility is established, several factors contribute to non-patching behavior: practitioners may be unaware of the system, unable to patch it, or the system may be in the process of being retired. In some cases, shutting down a system is considered an alternative to patching. These factors operate in an environment where municipalities must balance multiple dependencies—knowledge about vulnerabilities, awareness of which systems require patching, reliance on external IT providers, and the time needed to test and roll out patches. This already assumes patching is prioritized, which is often not the case. Practitioners noted that even these routine processes are difficult to organize.

In Chapter 3, we found that advisories issued by the national or sector CSIRT play an important role in convincing internal stakeholders to prioritize patching. These findings show that municipalities operate under tight resource constraints for both patching priorities and procedures. Our findings also suggest that the effectiveness of national vulnerability-response frameworks is not limited by the detection speed of vulnerabilities. Instead, mandated organizations such as the national CSIRT fail to adequately support constituents in providing all notifications. Nonetheless, municipalities continue to struggle with procedural frictions and organizational dependencies.

This situation is problematic, as the inability to patch or delays in patching significantly increase the risk of compromise. Practitioners may be aware of vulnerable systems yet still unable to patch them, either because of mandate issues or because their infrastructure is outsourced. As illustrated in Chapters 2 and 3, such outsourcing arrangements make municipalities dependent on third-party providers, which can stall remediation and extend exposure windows. This is alarming because many exploitations of critical vulnerabilities occur shortly after disclosure [192]. Moreover, the rapid growth in publicly reported vulnerabilities—18,323 CVEs in 2020, 40,303 in 2024, and 37,223 in 2025 [50]—exacerbates the problem: even timely patching now requires capacity beyond what many municipalities can sustain.

The combination of an increasing surface exposure, impeded patching, limited resources, and the expanding number of vulnerabilities keeps municipalities in a state of constant risk. The probabilities for compromise in the causal model, which is depicted again at the start of this section, in Figure 5.1 are therefore likely to rise. Consequently, the factor “Reactive Measures” will grow in importance. Incident response is thus likely to become more frequent in municipal environments, elevating the role of sector CSIRTs in providing support. However, Chapter 3 finds that municipal expectations of CSIRT involvement during incidents often diverge from what CSIRTs are mandated or resourced to provide. These misaligned expectations undermine trust, increase costs, and delay recovery. On top of that, municipalities are at risk due to a broken notification pipeline between the national CSIRT and its constituents. In short, these factors jeopardize the effectiveness of reactive security measures and amplify the harm of cyber incidents.

Under the NIS2 regulation, municipalities are expected to manage their own cyber-security. However, as our findings show, limited resources and a lack of expertise make it

unlikely that many can maintain an adequate security posture independently. It remains unclear how these constraints translate into the societal impact of incidents. However, given the public nature of municipalities and the public services they provide, it is highly likely that cyber incidents will undermine public trust and the stability of local governments.

Fourth, it is relatively easy to find vulnerable systems with publicly available tools, techniques, and services. In Chapter 2, vulnerable hosts are observed with publicly available services Censys and Shodan. While attribution of those systems is difficult, attackers have the advantage here. Particularly for targeted attackers, such as APTs, even with a single system correctly attributed to a municipality, attackers can quickly map additional municipal infrastructure using automated discovery techniques, including Certificate Transparency Logs, passive DNS, SPF record analysis, certificate mapping, reverse DNS, WHOIS lookups, and more. Given that municipalities are attractive targets because they contain significant personal data and provide critical public services that may not be interrupted, many criminal actors may also be interested in penetrating municipal systems. For example, extortion or blackmail for financial gain.

5

The automated discovery techniques will likely provide attackers with additional opportunities to find a vulnerability and break into the municipality. Concerningly, using many of these techniques, including passive scanning services like Censys and Shodan, doesn't give defenders a clue on who is targeting them: there is no attacker signal. This is where the value of attributed IOCs, examined in Chapter 4, may come in handy: attributed IOCs may reveal which actors are targeting a municipality. Consequently, this may help the network defenders of a municipality tailor their already limited resources for an appropriate defence. However, as shown in Chapter 4, the painted picture may turn out to be wrong. Depending on which threat intelligence vendor a municipality is using, it may get different actor attributions. This seems hardly appropriate for the actor-centric defense that is nowadays advocated as a best practice. Additionally, neither a single feed or combination of feeds provides a complete or accurate view of their threat landscape. The IOC feeds of commercial TI vendors, therefore, do not fully illuminate the "Threat" component from the socio-technical causal model in Figure 5.1. This prevents municipalities to deploy tailored security measures for specific actors targeting their organizations. This is problematic for municipalities, and many organizations, which manage scarce cyber resources. Moreover, the commercial IOCs are very expensive, and there is no full coverage, suggesting a municipality needs at least several expensive feeds to get a grasp of their threat landscape. Together, these results question the practicality of actor-centric defense for low-resource public organizations. The empirical inconsistencies across commercial vendors show that the attribution itself is unstable, thus, policies assuming stable actor knowledge are misaligned with reality.

Finally, Chapters 2 and 3 illustrate that feedback loops and learning mechanisms are vital yet largely missing. In the case of vulnerabilities, we observed one-way information flows: notifications are sent out, but few mechanisms exist to confirm remediation, assess outcomes, or learn from incidents. This absence prevents both municipalities and coordinating bodies from evaluating the effectiveness of their actions or improving processes over time.

The failures in the notification process reveal a systemic coordination problem among

actors in the cyber response ecosystem—the national CSIRT, the sector CSIRT, and the municipal constituents. These organizations share the same objective of reducing exposure, yet they lack reliable feedback and coordination mechanisms to verify progress toward that goal. Without these loops, each actor operates on partial information: the national CSIRT does not know whether alerts are received or acted upon, sector CSIRTs cannot measure their service effectiveness, and municipalities receive data without understanding how it fits into the broader response structure. In Chapter 3, the lack of feedback appears to underlie what we referred to as the “bootstrap problem” of sector CSIRTs. Without valuable notifications, municipalities have little incentive to update their asset data; yet without accurate asset data, the notifications lose value.

These findings suggest that the absence of feedback is not simply an issue of operational oversight but a structural weakness of the municipal cybersecurity ecosystem. Information moves efficiently in one direction but rarely returns to inform the next cycle of decision-making. Establishing formal feedback mechanisms, such as confirmation of receipt, remediation reporting, or periodic reviews, could therefore transform the notification process from a broadcast model into a learning system. Establishing effective feedback loops is essential to enabling a shift from simple information sharing to collective improvement at the municipal, sectoral, and national levels.

In sum, the findings across the three studies advance a more systemic understanding of municipal cybersecurity. They show that the effectiveness of security measures depends less on technology itself and more on how institutions organize, interpret, and act on security information. Improving municipal resilience, therefore, requires strengthening the links between detection, responsibility, and learning, rather than adding new tools. The studies reveal where coordination breaks down in practice and provide an empirical basis for designing security governance that better supports municipalities’ realities and constraints.

### 5.3. GOVERNANCE IMPLICATIONS

The findings across the three studies show that the effectiveness of security measures depends less on technologies and more on the institutional mechanisms that determine who acts, how information circulates, and how learning occurs. The findings can best be understood as outcomes of governance arrangements: the processes in which authority, incentives, information, and responsibility are structured across various interdependent actors in the cybersecurity domain [234].

A central characteristic of this setting is that many security failures generate system-level risk, i.e., failures at one organization impose costs on others or degrade the effectiveness of shared infrastructures and services, commonly known as negative externalities [10]. In the context of cybersecurity, externalities arise when the consequences of security decisions extend beyond the organization making them, while the costs of prevention and remediation are borne by organizations themselves. When such externalities accumulate across interconnected actors, such as municipalities, shared service providers, CSIRTs, or security vendors, then local underinvestment or coordination failure can undermine collective security outcomes.

To understand these challenges, this section adopts a governance perspective that distinguishes between three ‘ideal types’ or modes of coordination: *hierarchy*, *market*,

and *network* [125, 195, 241]. The modes each differ in how coordination is achieved. Hierarchies rely on authority, rules, and accountability; markets coordinate through prices, contracts, and competition; and networks depend on trust, reciprocity, and repeated interactions. None of these modes is inherently superior. As emphasized in [195], each has characteristic strengths and failure modes, and real-world governance arrangements are typically hybrid. Furthermore, governance outcomes depend less on the presence of any single mode and more on how modes are combined and aligned with the specific coordination problem [241]. Consequently, institutional arrangements shape what kinds of security outcomes are feasible, stable, and scalable [125].

Using the governance modes as an analytical lens, this section revisits the empirical findings of Chapters 2–4 and articulates implications for cybersecurity governance.

### 5.3.1. HIERARCHY

Hierarchical governance becomes salient where coordination problems involve system-level risk, unclear accountability, and persistent failures of voluntary coordination. Hierarchical governance is comparatively strong in producing reliability and accountability, but comes at the cost of flexibility and adaptability [195]. As Chapters 2 and 3 make clear, hierarchical governance shortcomings for municipalities in the cyber ecosystem are not primarily about excessive control, but rather, about insufficiently specified authority and responsibility.

First, the vulnerability-notification process documented in Chapter 3 and, to some extent, the diffuse responsibility for vulnerable hosts in Chapter 2 illustrate this gap. We identified a failure in the notification mechanism: large numbers of vulnerability notifications never reached constituents, and the national CSIRT did not detect the omissions. Notifications traverse from national CSIRTs to sector-level CSIRTs and ultimately to municipalities, yet no actor has a clearly defined obligation or authority to ensure end-to-end delivery. Furthermore, when notifications fail to propagate, this failure is difficult to detect and correct, and responsibility is very diffuse. Accountability is fragmented across organizational layers, without an accompanying mechanism for escalation or verification. The implication is not a need for tighter operational control, but instead for clarity about who is responsible for the integrity of the notification chain as a whole. For example, by specifying which actor is responsible for monitoring notification completeness, by defining escalation points when forwarding fails, or by requiring minimal reporting on notification coverage and timeliness. Such mechanisms aim to improve coordination reliability rather than rely solely on technical operational controls.

This lack of hierarchical clarity also contributes directly to the observed bootstrapping failure in asset registration. Municipalities that do not receive vulnerability notifications derive little immediate benefit from maintaining accurate asset data, while incomplete registrations further reduce the effectiveness of notification services. It thus seems that, in the absence of an empowered actor to intervene when participation declines or data quality degrades, the mechanism becomes locked in a downward spiral. Hierarchical governance may assign responsibility for sustaining the preconditions of coordination, such as minimum participation levels or data-quality expectations, periodic validation of registrations, or a designated authority mandated to intervene when data quality deteriorates. Here too, the mechanisms are not aimed at technical security

measures, but at governing participation and upkeep.

Closely related, top-down regulation could help turn asset management into a mature governance function for determining whether technical controls are functioning correctly. This might include mandating clear ownership for every system within an organization, including externally hosted and SaaS Solutions. As a result, this should bring the integrity of asset inventory into parity with financial reports and privacy compliance, as a key component of organizational integrity rather than an operational detail of the IT department. Chapters 2 and 3 underscore that many of the problems in vulnerability management and notification stem from outdated or incomplete asset data. Without accurate visibility into which systems belong to which organization, even the most sophisticated detection or notification efforts fail. Regulation may help national and sector CSIRTs, and their constituents, develop continuously updated asset and incident registries, ideally integrating automated discovery or EASM solutions. This would result in shared registries that could create a common operational picture and provide an empirical basis for measuring security progress across sectors.

Next, the outsourcing process and shared-service arrangements show how hierarchical governance can be misaligned. Chapters 2 and 3 showed that municipalities frequently delegate operational control over ICT systems while retaining political and legal accountability for security incidents. This separation of authority and responsibility weakens the municipality's ability to act on security information and complicates remediation. This represents a typical hierarchical coordination problem: those accountable for outcomes lack the authority to influence relevant decisions [241]. The implication is that new hierarchical mechanisms are needed to realign authority with accountability across service arrangements spanning multiple organizations. For example, authority can be realigned to accountability through enforceable reporting obligations from service providers, explicitly encoded patching procedures, clearly defined escalation rights for municipalities, or audit provisions that assess whether service arrangements allow accountable actors to exercise meaningful control.

One effect of hierarchical governance measures is increased cybersecurity requirements for municipalities. These requirements aim to improve security, but in reality, cybersecurity resource allocation is already lagging. For example, policies such as the EU's NIS2 Directive, itself a good example of hierarchical governance by mandating organizations with certain obligations, expect municipalities to be responsible for their own cybersecurity, yet the empirical findings suggest that municipalities lack the resources, expertise, and institutional capacity to do so effectively. Likewise, sector CSIRTs remain underfunded. Therefore, hierarchical governance might also focus on increasing efficiency and effectiveness. For example, to increase capabilities that focus on shared capacity-building, e.g., joint technical infrastructure or pooled security staff.

Furthermore, Chapter 3 revealed a mismatch in how the roles of sector CSIRTs are understood by various stakeholders. These kinds of mismatches derive from a lack of clear mandates and authority. After all, service implementation is context-specific for every sector CSIRT [74]. Here, regulation mechanisms may prove effective. Constituents often expected operational response and direct technical support, whereas CSIRTs conceive of their role as advisory and coordinative. This misalignment erodes trust and impedes collaboration. Regulation could help governance frameworks, or CSIRT commu-

nications, at a minimum, formalize service boundaries and escalation paths. Explicitly delineating which services are provided and under what conditions will help align resource allocation with expectations, improving both perceived legitimacy and actual effectiveness.

Finally, commercial security vendors, particularly in the area of threat intelligence, are often treated as authoritative. Yet, our empirical comparison reveals disagreement and incomplete coverage in threat actor reporting and IoC attribution. In the hierarchical mode, this suggests a mechanism of informal hierarchy based on reputation: without an authoritative actor, some actors are treated as such because of their perceived expertise [241]. This dynamic could prove dangerous. When the outputs of those vendors, specifically attributed IoCs, inform decision-making, reliance on reputational authority alone risks institutionalizing the inconsistencies observed in attributions. This suggests a role for hierarchical standard-setting, not to centralize the process or the act of attribution, but instead to explicitly articulate shared expectations regarding standards of evidence or confidence reporting.

## 5

### 5.3.2. MARKET

Market governance coordinates action through prices, contracts, and competition. It is most effective where quality is observable, goods are comparable, and costs and benefits accrue to the same actors. In the cybersecurity domain, though, many activities are characterized by information asymmetry, delayed or diffuse benefits, and limited opportunities for comparison [10]. For example, buyers of security services or threat intelligence often cannot directly assess quality at the time of procurement, and instead, must rely on reputational signals rather than verifiable performance [25]. Moreover, the benefits of security investments often materialize indirectly or accrue to actors outside the buyer–supplier relationship, leading to underinvestment in security efforts. Therefore, market-based coordination, i.e., market mechanisms, frequently fail to reliably align incentives and information in ways that produce effective collective security outcomes. Particularly for security services that are difficult to evaluate up front.

The difficulty in evaluating the quality of security services became clear in Chapter 4, where we found inconsistent IoC attribution in commercial threat intelligence. The empirical comparison of vendors shows substantial disagreement and limited overlap in attributed actors, even among widely used and costly services. While vendors operate in a competitive market, competition has not led to convergence or interoperability in attribution practices. Instead, differentiation in coverage, naming schemes, and analytic judgments appears to be a rational market strategy. From a governance perspective, this indicates that market coordination alone does not produce shared situational awareness when the quality of the service is inherently difficult to observe or verify.

A concrete implication is that procurement-based governance mechanisms cannot assume that buying more or different intelligence feeds will lead to convergence in how threats are identified and understood. Where attribution outputs from commercial vendors are used as inputs to coordinate or prioritize security efforts, such as during incident response or in tailoring organizational defences, market mechanisms may need to be complemented by governance arrangements that improve comparability. For example, there could be regulations on confidence levels or certain standards of evidence

when procuring security services. Or, vendors could be guided towards uniform actor naming schemes to improve cross-vendor comparisons. However, initiatives such as the threat actor naming deconfliction effort of Microsoft and CrowdStrike might be an early sign of the positive effect of market mechanisms [154]. Nonetheless, the proposed mechanisms do not regulate attribution outcomes but instead govern how attribution products can be meaningfully evaluated and combined across vendors.

Market limitations also become clear in municipal outsourcing and shared-service arrangements. Municipalities frequently procure ICT services in competitive markets, yet our findings show that such arrangements often undermine the municipality's ability to act on security information. As a result, price competition and contractual service levels do not reliably align supplier incentives with the municipality's risk exposure.

In sum, the findings suggest that market governance in municipal cybersecurity is insufficient when relied upon to ensure coordination, learning, or consistency across organizational boundaries. Without complementary mechanisms that address information asymmetry and externalities, market arrangements tend to foster fragmentation rather than resolve it.

### 5.3.3. NETWORK

Network governance coordinates action through trust, reciprocity, and repeated interaction rather than through authority or price. Networks are particularly effective where coordination depends on the exchange of rich, context-dependent information and where formal contracting or hierarchical control is impractical [195]. This mode is particularly suited for domains where information flows freely and frequently, and where there is little top-down regulation. Unsurprisingly, in Chapter 3 we found that many sector CSIRTS originated from bottom-up initiatives while top-down regulation was absent. Networks are therefore a natural governance mode for activities such as information sharing, incident coordination, and mutual assistance, where effectiveness depends on cooperation among interdependent actors rather than compliance with predefined rules.

The findings from Chapter 3 show that many core cybersecurity services operate de facto as networked arrangements. The effectiveness of vulnerability notification, incident coordination, and advisory services depends heavily on informal relationships, personal contacts, and mutual trust between CSIRT staff and municipal stakeholders. These relational ties often compensate for gaps in formal authority or contractual clarity and enable coordination to occur at all. It illustrates a key strength of network governance: it enables coordination even in the absence of clearly specified roles or enforceable obligations.

However, the findings also reveal characteristic failure modes of network governance. For example, the absence of systematic feedback mechanisms undermines reciprocity. As shown in Chapter 3, vulnerability notifications are typically sent downstream, but confirmation, remediation status, or learning rarely flows back upstream. Over time, this one-way exchange weakens trust and reduces participants' incentives to invest in shared coordination infrastructure, such as maintaining accurate asset registrations. From a network perspective, this is not merely an operational omission but a structural weakness: networks rely on reciprocal exchange to sustain cooperation, and when reciprocity erodes, coordination degrades. Examples of network-maintenance mechanisms include

structured feedback loops, notification audits, and shared asset registries that stabilize coordination beyond individual relationships. Such mechanisms are currently missing, as observed in Chapter 3.

The bootstrapping failure observed in poor asset registration further illustrates the limits of network governance at scale. Trust-based coordination may work within small or stable groups, but Chapter 3 showed that the municipal ecosystem is heterogeneous and dynamic, and municipalities vary greatly in capabilities and resources. In such settings, reliance on informal norms and personal relationships alone is insufficient to sustain consistent participation over time. The implication is that network governance requires explicit support structures such as routinized feedback practices, shared expectations about responsiveness, or clearly defined liaison roles, to remain viable as the number and diversity of participants increase.

## 5.4. FUTURE WORK

In this chapter, we look at avenues for future research. Chapters 2- 4 each address suggestions for future research, and therefore, we look at broader research directions that emerged from our work.

5

### MULTI-DISCIPLINARY RESEARCH INTO ASSET MANAGEMENT AND VULNERABILITY NOTIFICATIONS

Chapters 2 and 3 showed that asset management for municipalities is problematic. Assets are fleeting or unreported. One strand of future research should examine factors that affect or improve asset management. Do some organizations have features that support better asset management practices? And, what factors occur across various organizational types? Closely related, another research direction concerns the evaluation of the accuracy of techniques for mapping organizational assets. How well do public reconnaissance techniques work to identify an organization's assets? Given the growing importance of asset inventory, how well do available tools work? An interdisciplinary scientific study comparing the results of External Attack Surface Management tools with organizational ground-truth may uncover insights that bridge the organizational challenges this dissertation identified. Another strand of future work relates to vulnerability notification and patching. In this dissertation, we look at vulnerable systems identified via the CPE and known CVEs. However, many vulnerabilities are in systems that do not reveal a service or version in their banners. For example, various incidents have shown that many security products, such as firewalls and VPN servers, contain vulnerabilities. These vulnerabilities generally have a high severity rating. How well do practitioners patch such vulnerabilities, and what factors impact patching and non-patching behavior? Moreover, are such factors different across various organizational types? Or even across different types of vulnerabilities? Such studies provide additional, granular insight into patching behaviors that could greatly benefit policymakers. This line of research could also be extended to vulnerability notifications. Future work could use a somewhat similar research design to the work in Chapter 2, where vulnerable systems are measured, notifications are sent, and patching behaviors are measured, with complementary interviews that capture and contextualize the considerations made. Such studies could be repeated across different organizational types and industries. Those in-

sights would provide policymakers with actionable results to tailor advisories to specific constituents. With additional sectors under NIS2, many new organizations enter the cyber domain with cybersecurity obligations. These organizations are likely to operate in their own dynamics, which we have yet to uncover.

#### MEASURING THE EFFECTS AND MITIGATIONS AFTER COMPROMISE

This dissertation showed that vulnerable hosts are challenging to remediate. Yet vulnerabilities also emerge quickly. Consequently, organizations should have preventive security measures in place, but prepare for compromise. We currently lack empirical insights into how incident handling, mitigation, and recovery occur across organizations. What is the impact of compromise on reactive security measures and actual harm? What is the effect of the misaligned expectations that we observed in Chapter 3? New studies will face challenging measurement issues, but a combination of qualitative and quantitative methods will provide insights that are highly valuable to effective incident response, making cyber ecosystems much more resilient.

#### CROSS-COUNTRY MUNICIPAL, INSTITUTIONAL AND POLICY COMPARISONS

Chapters 2 and 3 provided case-studies of institutional structures within the Netherlands. The findings are likely to generalize to other countries with similar structures, but future work may look into case studies of different institutional structures. Or, institutional structures and their impact may be compared across countries. Academic studies on municipalities and security are limited, and the scientific community still has much to learn about how these organizations operate across different institutional structures worldwide.

#### EXPLORE CSIRT SERVICES

Academic work on Computer Security Incident Response (CSIRT) teams is very limited. In Chapter 3, we shed some light on the challenges faced by sector CSIRTs, but much remains to be uncovered. First, CSIRTs come in various shapes, sizes, and forms, as exemplified by the team types identified in the FIRST framework [77]. Future work may look into the specifics of each of these team types. Where are they different? Under what conditions do they operate optimally? Does that differ by country or by sector? Such studies would provide insights into the quality evaluation or performance evaluation of these organization types that we currently lack. Second, we conducted an exploratory study of the service challenges faced by sector CSIRTs. In doing so, we focused on the vulnerability notification mechanism. However, new studies may focus on other services with more specific measurement methodologies. Those studies might reveal the conditions under which CSIRT services might function correctly. Third, this dissertation advocates the use of feedback mechanisms. However, we lack empirical data on existing feedback mechanisms and their quality evaluation. What feedback mechanisms exist, how do they work, and when do they work best? Answers to those questions will help policymakers to effectively design top-down regulations that make the cyber ecosystem more resilient.

#### ATTRIBUTION IN ATTACKER STUDIES

In Chapter 4 we compared IOC attribution across seven commercial vendors. This study revealed that actor attribution is not stable. Future work may investigate the root causes of this instability. What difficulties do analysts in this domain face? What factors may have led to different attributions? In our study, we focused only on commercial sources, but research gaps remain in systematic insights into attribution across public sources, as well as between public and closed (i.e., commercial) sources and vendors. Results may provide valuable insights for organizations to help determine whether pricy commercial TI is worth the investment, or under what conditions it may be. Furthermore, many existing studies on attackers and TIs lack ground truth data. Access to ground truth data would be very valuable for assessing the quality of both public and commercial TI services and for revealing more about attacker modus operandi. Academics should explore collaborations with organizations that have access to ground-truth data, such as police organizations. Analyses of seized datasets by academics may lead to novel scientific insights that benefit both the scientific and law enforcement communities.

# BIBLIOGRAPHY

- [1] Maziana Abd Majid and Khairul Akram Zainol Ariffin. 2021. Model for successful development and implementation of Cyber Security Operations Centre (SOC). *PloS One* 16, 11 (2021), e0260157. <https://doi.org/10.1371/journal.pone.0260157>
- [2] Lawrence Abrams. 2020. Citrix ADC CVE-2019-19781 Exploits Released, Fix Now! Online: <https://www.bleepingcomputer.com/news/security/citrix-adc-cve-2019-19781-exploits-released-fix-now/>. <https://www.bleepingcomputer.com/news/security/citrix-adc-cve-2019-19781-exploits-released-fix-now/> Accessed: 2023-06-01.
- [3] AfricaCERT. 2023. Mission statement. Online: <https://www.africacert.org/mission-statement/>. <https://www.africacert.org/mission-statement/> Accessed: 2023-09-6.
- [4] Gabriela Ahmadi-Assalemi, Haider Al-Khateeb, Gregory Epiphaniou, and Carsten Maple. 2020. Cyber resilience and incident response in smart cities: A systematic literature review. *Smart Cities* 3, 3 (2020), 894–927.
- [5] Bushra A. Alahmadi, Louise Axon, and Ivan Martinovic. 2022. 99% False Positives: A Qualitative Study of SOC Analysts' Perspectives on Security Alarms. In *31st USENIX Security Symposium (USENIX Security 22)*. USENIX Association, Boston, MA, 2783–2800. <https://www.usenix.org/conference/usenixsecurity22/presentation/alahmadi>
- [6] Luca Allodi, Fabio Massacci, and Julian Williams. 2022. The Work-Averse Cyberattacker Model: Theory and Evidence from Two Million Attack Signatures. *Risk Analysis* 42, 8 (2022), 1623–1642. <https://doi.org/10.1111/risa.13732> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/risa.13732>
- [7] Noura Alomar, Primal Wijesekera, Edward Qiu, and Serge Egelman. 2020. You've Got Your Nice List of Bugs, Now What? Vulnerability Discovery and Management Processes in the Wild. In *Proceedings of the Sixteenth Symposium on Usable Privacy and Security (SOUPS 2020)*. USENIX Association, Boston, MA, USA, 319–339. <https://www.usenix.org/conference/soups2020/presentation/alomar>
- [8] Noura Alomar, Primal Wijesekera, Edward Qiu, and Serge Egelman. 2020. "You've Got Your Nice List of Bugs, Now What?" Vulnerability Discovery and Management Processes in the Wild. In *Sixteenth Symposium on Usable Privacy and Security (SOUPS 2020)*. USENIX Association, Berkeley, CA, USA, 319–339. <https://www.usenix.org/conference/soups2020/presentation/alomar>

- [9] Florian J. Egloff and. 2020. Contested public attributions of cyber incidents and the role of academia. *Contemporary Security Policy* 41, 1 (2020), 55–81. <https://doi.org/10.1080/13523260.2019.1677324> arXiv:<https://doi.org/10.1080/13523260.2019.1677324>
- [10] Ross Anderson and Tyler Moore. 2006. The Economics of Information Security. *Science* 314, 5799 (2006), 610–613. <https://doi.org/10.1126/science.1130992>
- [11] Annika Andreasson, Henrik Artman, Joel Brynielsson, and Ulrik Franke. 2024. Cybersecurity work at Swedish administrative authorities: taking action or waiting for approval. *Cognition, Technology & Work* 26, 4 (2024), 709–731.
- [12] Atlas.Ti. 2023. ATLAS.ti | The #1 Software for Qualitative Data Analysis. Online: <https://atlasti.com>. <https://atlasti.com> Accessed: 2023-03-01.
- [13] M Bada, S Creese, M Goldsmith, C Mitchell, and E Phillips. 2014. Improving the Effectiveness of CSIRTs Global Cyber Security Capacity Centre.
- [14] Priyanka Badva, Kopo M. Ramokapane, Eleonora Pantano, and Awais Rashid. 2024. Unveiling the Hunter-Gatherers: Exploring Threat Hunting Practices and Challenges in Cyber Defense. In *33rd USENIX Security Symposium (USENIX Security 24)*. USENIX Association, Philadelphia, PA, 3313–3330. <https://www.usenix.org/conference/usenixsecurity24/presentation/badva>
- [15] European Central Bank. 2025. TIBER-EU Targeted Threat Intelligence Report Guidance.
- [16] Gil Baram. 2024. Cyber Diplomacy through Official Public Attribution: Paving the Way for Global Norms. *International Studies Perspectives* advance online publication, advance online publication (2024), ekae022. <https://doi.org/10.1093/isp/ekae022> arXiv:<https://academic.oup.com/isp/advance-article-pdf/doi/10.1093/isp/ekae022/60734708/ekae022.pdf>
- [17] Sean Barnum. 2012. Standardizing Cyber Threat Intelligence Information with the Structured Threat Information Expression (stix). *MITRE Corporation* 11, 2012 (2012), 1–22.
- [18] Frederick Barr-Smith, Xabier Ugarte-Pedrero, Mariano Graziano, Riccardo Spolaor, and Ivan Martinovic. 2021. Survivalism: Systematic Analysis of Windows Malware Living-Off-The-Land. In *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, Los Alamitos, CA, USA, 1557–1574. <https://doi.org/10.1109/SP40001.2021.00047>
- [19] Frederick Barr-Smith, Xabier Ugarte-Pedrero, Mariano Graziano, Riccardo Spolaor, and Ivan Martinovic. 2021. Survivalism: Systematic Analysis of Windows Malware Living-Off-The-Land. *IEEE Symposium on Security and Privacy* 2021, 1 (2021), 1557–1574.

- [20] Rob Barrett, Eser Kandogan, Paul P. Maglio, Eben M. Haber, Leila A. Takayama, and Madhu Prabhaker. 2004. Field studies of computer system administrators: analysis of system management tools and practices. In *Proceedings of the 2004 ACM conference on Computer supported cooperative work (CSCW '04)*. Association for Computing Machinery, New York, NY, USA, 388–395. <https://doi.org/10.1145/1031607.1031672>
- [21] Adam Beautement, M. Angela Sasse, and Mike Wonham. 2008. The compliance budget: managing security behaviour in organisations. In *Proceedings of the 2008 New Security Paradigms Workshop (NSPW '08)*. Association for Computing Machinery, New York, NY, USA, 47–58. <https://doi.org/10.1145/1595676.1595684>
- [22] H. L. J. Bijmans and M. S. C. van Leuken. 2024. No Time to Choose: Leveraging Internet Scans to Determine IoC Lifetimes. In *2024 IEEE International Conference on Big Data (BigData)*. IEEE, Los Alamitos, CA, USA, 2586–2595. <https://doi.org/10.1109/BigData62323.2024.10825640>
- [23] Peter Block. 2018. *Community: The structure of belonging*. Berrett-Koehler Publishers, Oakland, CA, USA.
- [24] Tamara G. Bondar, Hala Assal, and Abdelrahman Abdou. 2023. Why do Internet Devices Remain Vulnerable? A Survey with System Administrators. In *Proceedings of the Workshop on Measurements, Attacks, and Defenses for the Web (MADWeb 2023)*. Internet Society, San Diego, CA, USA, n/a. <https://www.ndss-symposium.org/ndss-program/madweb-2023/>
- [25] Xander Bouwman, Harm Griffioen, Jelle Egbers, Christian Doerr, Bram Klievink, and Michel van Eeten. 2020. A different cup of TI? The added value of commercial threat intelligence. In *29th USENIX Security Symposium (USENIX Security 20)*. USENIX Association, Online, 433–450. <https://www.usenix.org/conference/usenixsecurity20/presentation/bouwman>
- [26] Virginia Braun and Victoria Clarke. 2021. One size fits all? What counts as quality practice in (reflexive) thematic analysis? *Qualitative Research in Psychology* 18, 3 (2021), 328–352. <https://doi.org/10.1080/14780887.2020.1769238> arXiv:<https://doi.org/10.1080/14780887.2020.1769238>
- [27] Bughra. 2025. Weaponization Techniques for Red Team Operations. <https://bughra.dev/posts/weaponization/> Accessed: 2025-08-26.
- [28] Sergio Caltagirone, Andrew Pendergast, and Christopher Betz. 2013. The diamond model of intrusion analysis. *Threat Connect* 298, 0704 (2013), 1–61.
- [29] Censys. 2023. Frequently Asked Questions. Online: <https://support.censys.io/hc/en-us/articles/360038378552-Frequently-Asked-Questions>. <https://support.censys.io/hc/en-us/articles/360038378552-Frequently-Asked-Questions> Accessed: 2023-03-2.

- [30] Center for Internet Security. 2023. 2022 NCSR: SLTTs Excel in Recovery Planning and Mitigation. <https://www.cisecurity.org/insights/blog/2022-ncsr-slttps-excel-in-recovery-planning-and-mitigation> Accessed: 8-10-2025.
- [31] UK National Cyber Security Centre. 2025. Cyber Essentials – Overview. <https://www.ncsc.gov.uk/cyberessentials/overview>. Accessed: 29-9-2025.
- [32] Thai CERT. 2025. All groups - Threat Group Cards: A Threat Actor Encyclopedia. <https://apt.etcha.or.th/cgi-bin/listgroups.cgi> Accessed: 2025-07-30.
- [33] CERT.br. 2023. About CERT.br. Online: <https://www.cert.br/about/>. <https://www.cert.br/about/> Accessed: 2023-09-6.
- [34] F. Cetin, C. Gañán, Maciej Korczyński, and M. van Eeten. 2017. Make Notifications Great Again: Learning How to Notify in the Age of Large-Scale Vulnerability Scanning. Workshop on the Economics of Information Security (WEIS). <https://www.semanticscholar.org/paper/Make-notifications-great-again%3A-learning-how-to-in-Cetin-Ga%C3%B1%C3%A1n/ed24ca9d63385392bbd6ac52288933b93444c43d> Paper.
- [35] Tiffani R. Chen, Daniel B. Shore, Stephen J. Zaccaro, Reeshad S. Dalal, Lois E. Tetric, and Aiva K. Gorab. 2014. An Organizational Psychology Perspective to Examining Computer Security Incident Response Teams. *IEEE Security & Privacy* 12, 5 (Sept. 2014), 61–67. <https://doi.org/10.1109/MSP.2014.85> Conference Name: IEEE Security & Privacy.
- [36] Niraj Chokshi. 2019. Hackers Are Holding Baltimore Hostage: How They Struck and What’s Next. *The New York Times*. <https://www.nytimes.com/2019/05/22/us/baltimore-ransomware.html> Accessed: 2023-05-15.
- [37] Chris Teale (GCN). 2023. Unpatched, known vulnerabilities still key driver of cyberattacks. Online: <https://gcn.com/cybersecurity/2023/03/unpatched-known-vulnerabilities-still-key-driver-cyberattacks/383489/>. Accessed: 2023-05-15.
- [38] Justin Novak Christopher Rodman, Breanna Kraus. 2024. SOC Service Areas: Identification, Prioritization, and Implementation. <https://www.ndss-symposium.org/ndss-paper/auto-draft-521/>
- [39] CISA. 2023. *Exploitation of Unitronics PLCs used in Water and Wastewater Systems*. Alert. Cybersecurity and Infrastructure Security Agency (CISA). <https://www.cisa.gov/news-events/alerts/2023/11/28/exploitation-unitronics-plcs-used-water-and-wastewater-systems> Accessed: 1-10-2025.
- [40] CISA and FBI. 2020. *Russian State-Sponsored Advanced Persistent Threat Actor Compromises U.S. Government Targets*. Cybersecurity Advisory AA20-296A. U.S. Department of Homeland Security / CISA. <https://www.cisa.gov/news->

- events/cybersecurity-advisories/aa20-296a Last revised December 01, 2020.
- [41] City Auditor's Office City of Atlanta. 2018. *ISO/IEC 27001 ISMS Precertification Audit — January 2018*. Technical Report. City Auditor's Office, City of Atlanta. <https://www.atlaudit.org/isoiec-27001-isms-precertification-audit---january-2018.html> Accessed: 29-9-2025.
- [42] Department of Information & Technology Services City of Dallas. 2023. *The City of Dallas Ransomware Incident: May 2023 — After Action Review*. Technical Report. City of Dallas. <https://dallascityhall.com/departments/ciservices/DCH%20Documents/dallas-after-action-review-slideshow-of-ransomware-attack.pdf> Accessed: 30-9-2025.
- [43] European Commission. 2025. NIS2 Directive. <https://digital-strategy.ec.europa.eu/en/policies/nis2-directive>. Accessed: 29-9-2025.
- [44] European Commission. 2025. NIS2 Directive: new rules on cybersecurity of network and information systems | Shaping Europe's digital future. <https://digital-strategy.ec.europa.eu/en/policies/nis2-directive>
- [45] Luke Connolly. 2025. The State of Ransomware in the U.S.: Report and Statistics 2024. <https://www.emsisoft.com/en/blog/46288/the-state-of-ransomware-in-the-u-s-report-and-statistics-2024/> Accessed: 8-10-2025.
- [46] Joseph Cox. 2020. Inside Hackney's Struggle to Recover from a Ransomware Attack. <https://www.wired.com/story/ransomware-attack-recovery-hackney/>. Accessed: 29-9-2025.
- [47] Andrea Cristaldi. 2025. APTMap. <https://github.com/andreacristaldi/APTmap> original-date: 2020-07-27T06:06:36Z.
- [48] CrowdStrike. 2025. What is Cyber Threat Intelligence? [Beginner's Guide] - CrowdStrike. <https://www.crowdstrike.com/en-us/cybersecurity-101/threat-intelligence/> Accessed: 2025-07-07.
- [49] Garrett Cullity. 2025. The Free Rider Problem. In *The Stanford Encyclopedia of Philosophy* (Fall 2025 ed.), Edward N. Zalta and Uri Nodelman (Eds.). Metaphysics Research Lab, Stanford University, Stanford, CA, USA.
- [50] CVEdetails.com. 2025. CVE Vulnerabilities by Date. <https://www.cvedetails.com/browse-by-date.php>. Accessed: 13-10-2025.
- [51] CyberSaint. 2024. Recommendations for your next CIS risk assessment. <https://www.cybersaint.io/blog/recommendations-for-your-next-cis-risk-assessment> Accessed: 29-09-2025.

- [52] A. D'Amico and K. Whitley. 2008. The Real Work of Computer Network Defense Analysts. In *VizSEC 2007: Proceedings of the Workshop on Visualization for Computer Security*, John R. Goodall, Gregory Conti, and Kwan-Liu Ma (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 19–37. [https://doi.org/10.1007/978-3-540-78243-8\\_2](https://doi.org/10.1007/978-3-540-78243-8_2)
- [53] Davidjbianco. 2013. Enterprise Detection & Response: The Pyramid of Pain. <https://detect-respond.blogspot.com/2013/03/the-pyramid-of-pain.html>
- [54] Stephanie de Smale, Rik van Dijk, Xander Bouwman, Jeroen van der Ham, and Michel van Eeten. 2023. No One Drinks From the Firehose: How Organizations Filter and Prioritize Vulnerability Information. In *IEEE Symposium on Security and Privacy (SP) (SP)*. IEEE Computer Society, Los Alamitos, CA, USA, 203–219. <https://doi.org/10.1109/SP46215.2023.00012>
- [55] Stephen Deere. 2018. Confidential report: Atlanta's cyber attack could cost taxpayers \$17 million. <https://www.ajc.com/news/confidential-report-atlanta-cyber-attack-could-hit-million/GAljmndAF3EQdVWlMcXSOK/>. Accessed: 29-9-2025.
- [56] Constanze Dietrich, Katharina Krombholz, Kevin Borgolte, and Tobias Fiebig. 2018. Investigating System Operators' Perspective on Security Misconfigurations. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security (CCS '18)*. Association for Computing Machinery, New York, NY, USA, 1272–1289. <https://doi.org/10.1145/3243734.3243794>
- [57] Dulaunoy, Alexandre and Roth, Florian. 2025. Threat Actor - MISP galaxy. <https://misp-galaxy.org/threat-actor/> Accessed: 2025-06-10.
- [58] Zakir Durumeric, David Adrian, Ariana Mirian, Michael Bailey, and J. Alex Halderman. 2015. A Search Engine Backed by Internet-Wide Scanning. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security (CCS '15)*. Association for Computing Machinery, New York, NY, USA, 542–553. <https://doi.org/10.1145/2810103.2813703>
- [59] Zakir Durumeric, Frank Li, James Kasten, Johanna Amann, Jethro Beekman, Mathias Payer, Nicolas Weaver, David Adrian, Vern Paxson, Michael Bailey, and J. Alex Halderman. 2014. The Matter of Heartbleed. In *Proceedings of the 2014 Conference on Internet Measurement Conference (IMC '14)*. Association for Computing Machinery, New York, NY, USA, 475–488. <https://doi.org/10.1145/2663716.2663755>
- [60] Zakir Durumeric, Frank Li, James Kasten, Johanna Amann, Jethro Beekman, Mathias Payer, Nicolas Weaver, David Adrian, Vern Paxson, Michael Bailey, and J. Alex Halderman. 2014. The Matter of Heartbleed. In *Proceedings of the 2014 Conference on Internet Measurement Conference (IMC '14)*. Association for Computing Machinery, New York, NY, USA, 475–488. <https://doi.org/10.1145/2663716.2663755>

- [61] Zakir Durumeric, Eric Wustrow, and J. Alex Halderman. 2013. ZMap: Fast Internet-wide Scanning and Its Security Applications. In *22nd USENIX Security Symposium (USENIX Security 13)*. USENIX Association, Washington, DC, USA, 605–620. <https://www.usenix.org/conference/usenixsecurity13/technical-sessions/paper/durumeric>
- [62] Dutch Safety Board. 2019. Vulnerable through software - Lessons resulting from security breaches relating to Citrix software. Online: <http://www.onderzoeksraad.nl/en/page/17171/vulnerable-through-software---lessons-resulting-from-security>. <http://www.onderzoeksraad.nl/en/page/17171/vulnerable-through-software---lessons-resulting-from-security> Accessed: 2023-05-15.
- [63] Harun Ecik. 2021. Comparison of Active Vulnerability Scanning vs. Passive Vulnerability Detection. In *2021 International Conference on Information Security and Cryptology (ISCTURKEY)*. IEEE, Ankara, Turkey, 87–92.
- [64] ENISA. 2019. Study on CSIRT landscape and IR capabilities in Europe 2025. <https://www.enisa.europa.eu/publications/study-on-csirt-landscape-and-ir-capabilities-in-europe-2025>
- [65] ENISA. 2022. CSIRT Maturity Framework. <https://www.enisa.europa.eu/topics/incident-response/csirt-capabilities/csirt-maturity>. Accessed December 18, 2024.
- [66] ENISA. 2024. How to set up CSIRT and SOC | ENISA. <https://www.enisa.europa.eu/publications/how-to-set-up-csirt-and-soc>
- [67] ENISA. 2024. NIS Directive and national CSIRTs | ENISA. <https://www.enisa.europa.eu/publications/nis-directive-and-national-csirts>
- [68] ENISA. 2025. ENISA Threat Landscape 2024 | ENISA. <https://www.enisa.europa.eu/publications/enisa-threat-landscape-2024>
- [69] Aksel Ethembabaoglu, Natalia {I. Kadenko}, Yana Angelova, Yury Zhauniarovich, Rolf {van Wegberg}, Simon Parkin, and Michel {van Eeten}. 2026. “Tell Them They Are a Responsible Entity, Not a Customer”: Understanding Practitioner Challenges in Sector CSIRTs. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*, Nuria Oliver, {David A.} Shamma, and Heloisa Candello (Eds.). Association for Computing Machinery (ACM), United States, 1–23. <https://doi.org/10.1145/3772318.3790613> 2026 CHI Conference on Human Factors in Computing Systems, CHI 2026, CHI '26 ; Conference date: 13-04-2026 Through 17-04-2026.
- [70] Aksel Ethembabaoglu, Rolf van Wegberg, Yury Zhauniarovich, and Michel van Eeten. 2024. The Unpatchables: Why Municipalities Persist in Running Vulnerable Hosts. In *33rd USENIX Security Symposium (USENIX Security 24)*. USENIX Association, Philadelphia, PA, 7049–7066. <https://www.usenix.org/conference/usenixsecurity24/presentation/ethembabaogCollierlu>

- [71] European Union Agency for Cybersecurity (ENISA). 2025. ENISA Threat Landscape 2025. <https://doi.org/10.2824/1946> Accessed 2025-10-08.
- [72] John Everson and Peggy Cheng. 2024. A Survey on Network Attack Surface Mapping. *ACM Digital Threats: Research and Practice* 6, 1 (2024), 1–34. <https://doi.org/10.1145/3640019>
- [73] O. I. Falowo, K. Koshedo, and M. Ozer. 2023. An Assessment of Capabilities Required for Effective Cybersecurity Incident Management: A Systematic Literature Review. In *Proceedings of the 2023 IEEE International Conference on Digital Security and Privacy (DSPP)*. IEEE, Piscataway, NJ, USA, 1–11. <https://doi.org/10.1109/DSPP58763.2023.10404318>
- [74] FIRST. 2019. CSIRT Services Framework Version 2.1. Online: [https://www.first.org/standards/frameworks/csirts/csirt\\_services\\_framework\\_v2.1](https://www.first.org/standards/frameworks/csirts/csirt_services_framework_v2.1). [https://www.first.org/standards/frameworks/csirts/csirt\\_services\\_framework\\_v2.1](https://www.first.org/standards/frameworks/csirts/csirt_services_framework_v2.1) Accessed: 2023-05-15.
- [75] FIRST. 2019. CVSS v3.1 Specification Document. Online: <https://www.first.org/cvss/specification-document>. <https://www.first.org/cvss/specification-document> Accessed: 2023-05-15.
- [76] FIRST.ORG. 2019. FIRST CSIRT Services Framework. [https://www.first.org/standards/frameworks/csirts/FIRST\\_CSIRT\\_Services\\_Framework\\_v2.1.0\\_bugfix1.pdf](https://www.first.org/standards/frameworks/csirts/FIRST_CSIRT_Services_Framework_v2.1.0_bugfix1.pdf) Accessed: 2024-Nov-23.
- [77] FIRST.ORG. 2019. Team Types Within the Context of Services Frameworks. [https://www.first.org/standards/frameworks/csirts/team-type\\_1-0](https://www.first.org/standards/frameworks/csirts/team-type_1-0) Accessed: 2024-Nov-23.
- [78] US Office for Foreign Assets Control. 2021. Updated advisory on potential sanctions risks for facilitating ransomware payments. <https://ofac.treasury.gov/media/912981/download?inline> Accessed: 2025-07-07.
- [79] Federal Office for Information Security. 2023. CERT-Bund. Online: <https://www.bsi.bund.de/EN/Themen/Unternehmen-und-Organisationen/Cyber-Sicherheitslage/Reaktion/CERT-Bund/cert-bund.html?nn=907524>. <https://www.bsi.bund.de/EN/Themen/Unternehmen-und-Organisationen/Cyber-Sicherheitslage/Reaktion/CERT-Bund/cert-bund.html?nn=907524> Accessed: 2023-09-6.
- [80] Center for Internet Security. 2016. *CIS Input to the Commission on Cybersecurity RFI*. Technical Report. National Institute of Standards and Technology / U.S. Government. [https://www.nist.gov/system/files/documents/2016/09/15/cis\\_rfi\\_response.pdf](https://www.nist.gov/system/files/documents/2016/09/15/cis_rfi_response.pdf) Accessed: 29-9-2025.

- [81] Center for Internet Security. 2025. CIS Controls. <https://www.cisecurity.org/controls>. Accessed: 29-9-2025.
- [82] International Organization for Standardization. 2022. ISO/IEC 27001 — Information security, cybersecurity and privacy protection — Information security management systems — Requirements. <https://www.iso.org/standard/27001>. Accessed: 29-9-2025.
- [83] International Organization for Standardization. 2022. ISO/IEC 27002:2022 — Information security, cybersecurity and privacy protection — Information security controls. <https://www.iso.org/standard/75652.html>. Accessed: 29-9-2025.
- [84] Fox-SRT. 2022. CVE-2022-27510, CVE-2022-27518 – Measuring Citrix ADC & Gateway version adoption on the Internet. Online: <https://blog.fox-it.com/2022/12/28/cve-2022-27510-cve-2022-27518-measuring-citrix-adc-gateway-version-adoption-on-the-internet/>. Accessed: 2023-05-04.
- [85] Ashlee Frandell and Mary Feeney. 2022. Cybersecurity threats in local government: A sociotechnical perspective. *The American Review of Public Administration* 52, 8 (2022), 558–572.
- [86] Ashlee Frandell and Mary Feeney. 2022. Cybersecurity Threats in Local Government: A Sociotechnical Perspective. *The American Review of Public Administration* 52, 8 (2022), 558–572. <https://doi.org/10.1177/02750740221125432> arXiv:<https://doi.org/10.1177/02750740221125432>
- [87] Benjamin Freed. 2019. Baltimore approves \$10 million for ransomware recovery. <https://statescoop.com/baltimore-city-council-approves-10-million-ransomware-recovery/>. Accessed: 29-9-2025.
- [88] Konstantinos Fysarakis, Vasileios Mavroeidis, Manos Athanatos, George Spanoudakis, and Sotiris Ioannidis. 2022. A Blueprint for Collaborative Cybersecurity Operations Centres with Capacity for Shared Situational Awareness, Coordinated Response, and Joint Preparedness. In *Proceedings of the 2022 IEEE International Conference on Big Data (Big Data)*. IEEE, Piscataway, NJ, USA, 2601–2609. <https://doi.org/10.1109/BigData55660.2022.10020736>
- [89] Gartner. 2025. Best Security Threat Intelligence Products and Services Reviews 2025 | Gartner Peer Insights. <https://www.gartner.com/reviews/market/security-threat-intelligence-products-and-services> Accessed: 2025-08-12.
- [90] Kenneth Geers. 2010. The challenge of cyber attack deterrence. *Computer Law and Security Review* 26, 3 (2010), 298–303. <https://doi.org/10.1016/j.clsr.2010.03.003>
- [91] Béla Genge and Călin Enăchescu. 2016. ShoVAT: Shodan-based vulnerability assessment tool for Internet-facing services. *Security and Communication Networks* 9, 15 (2016), 2696–2714. <https://doi.org/10.1002/sec.1262> \_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/sec.1262>.

- [92] GitHub for MISP Threat Actor Galaxy. 2025. `misp-galaxy/clusters/threat-actor.json` at main · MISP/misp-galaxy. <https://github.com/MISP/misp-galaxy/blob/main/clusters/threat-actor.json> Accessed: 2025-07-30.
- [93] Jason Gray, Daniele Sgandurra, Lorenzo Cavallaro, and Jorge Blasco Alis. 2024. Identifying Authorship in Malicious Binaries: Features, Challenges & Datasets. *ACM Comput. Surv.* 56, 8, Article 212 (April 2024), 36 pages. <https://doi.org/10.1145/3653973>
- [94] Jonathan Greig. 2023. State of emergency declared as City of Oakland grapples with ransomware attack. Online: <https://therecord.media/oakland-ransomware-emergency-declared>. <https://therecord.media/oakland-ransomware-emergency-declared> Accessed: 2023-05-11.
- [95] Harm Griffioen, Tim Booiij, and Christian Doerr. 2020. Quality Evaluation of Cyber Threat Intelligence Feeds. In *Applied Cryptography and Network Security: 18th International Conference, ACNS 2020, Rome, Italy, October 19–22, 2020, Proceedings, Part II*. Springer-Verlag, Berlin, Heidelberg, 277–296. [https://doi.org/10.1007/978-3-030-57878-7\\_14](https://doi.org/10.1007/978-3-030-57878-7_14)
- [96] BOC Group. 2025. BSI IT Baseline Protection: an Overview. <https://www.boc-group.com/en/blog/grc/bsi-it-baseline-protection/>. Accessed: 29-9-2025.
- [97] Muhammad Haidar, Yudho Giri Suchahyo, Teddy Sukardi, and Arfive Gandhi. 2021. Analysis of CSIRT Services in Facing Cyber Security Challenges in Indonesia. In *Proceedings of the 4th International Conference on Information and Communications Technology (ICOIACT 2021)*. IEEE, Piscataway, NJ, USA, 154–159. <https://doi.org/10.1109/ICOIACT53268.2021.9563925>
- [98] Dean Hammer and Aaron Wildavsky. 1993. The Open-Ended, Semistructured Interview: An (Almost) Operational Guide. In *Craftways* (2 ed.). Routledge, London, UK.
- [99] Otto Hellwig, Gerald Quirchmayr, Edith Huber, Gernot Goluch, Franz Vock, and Bettina Pospisil. 2016. Major Challenges in Structuring and Institutionalizing CERT-Communication. In *Proceedings of the 11th International Conference on Availability, Reliability and Security (ARES 2016)*. IEEE, Piscataway, NJ, USA, 661–667. <https://doi.org/10.1109/ARES.2016.57>
- [100] Sk Tahsin Hossain, Tan Yigitcanlar, Kien Nguyen, and Yue Xu. 2024. Local Government Cybersecurity Landscape: A Systematic Review and Conceptual Framework. *Applied Sciences* 14, 13 (2024), 5501. <https://doi.org/10.3390/app14135501>
- [101] Allen Householder and Jonathan Spring. 2022. *A State-Based Model for Multi-Party Coordinated Vulnerability Disclosure (MPCVD)*. Technical Report. Carnegie Mellon University. <https://doi.org/10.1184/R1/16416771.v1>

- [102] IB&P. 2022. Kwetsbaar door software - Lessen naar aanleiding van beveiligingslekken door software van Citrix - Informatiebeveiliging & Privacy. Online: <https://ib-p.nl/download/kwetsbaar-door-software-lessen-naar-aanleiding-van-beveiligingslekken-door-software-van-citrix/>. Accessed: 2023-05-24.
- [103] Federal Bureau of Investigation (FBI) / IC3. 2021. *APT Actors Exploiting Fortinet Vulnerabilities to Gain Access for Malicious Activity*. FLASH / Advisory MI-000148-MW. FBI / IC3 / U.S. Department of Justice. <https://www.ic3.gov/CSA/2021/210527.pdf> Accessed: 1-10-2025.
- [104] Beyond Identity. 2025. All Threat Actors, APTs and Known Groups. <https://breach-hq.com/threat-actors> Accessed: 2025-07-30.
- [105] Ponemon Institute. 2018. The Value Of Threat Intelligence: The Second Annual Study Of North American & United Kingdom Companies. [https://library.cyentia.com/report/report\\_001935.html](https://library.cyentia.com/report/report_001935.html) Section: report.
- [106] SANS Institute. 2019. *The Evolution of Cyber Threat Intelligence (CTI): 2019 SANS CTI Survey | SANS Institute*. Technical Report. SANS Institute. <https://www.sans.org/white-papers/38790/> accessed: 2025-06-21.
- [107] Vincenzo Iozzo. 0100. The Case for Scale in Cyber Security. [https://media.ccc.de/v/36c3-11220-the\\_case\\_for\\_scale\\_in\\_cyber\\_security](https://media.ccc.de/v/36c3-11220-the_case_for_scale_in_cyber_security) Accessed: 2025-06-12.
- [108] Margaret Jack and Steven J. Jackson. 2016. Logistics as Care and Control: An Investigation into the UNICEF Supply Division. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 2209–2219.
- [109] Christopher Johnson, Mark Badger, David Waltermire, Julie Snyder, and Clem Skorupka. 2016. *Guide to Cyber Threat Information Sharing*. Technical Report NIST Special Publication (SP) 800-150. National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-150>
- [110] Jeffrey M. Jones. 2023. Americans Trust Local Government Most, Congress Least. <https://news.gallup.com/poll/512651/americans-trust-local-government-congress-least.aspx>. Accessed: 30-9-2025.
- [111] J. K. Kampen, S. Van de Walle, and G. Bouckaert. 2006. The impact of past performance on public trust: The role of service delivery in shaping attitudes toward government. *American Review of Public Administration* 36, 6 (2006), 576–594. <https://doi.org/10.1080/15309576.2006.11051881>
- [112] Georgia Killcrece, Klaus-Peter Kossakowski, Robin M. Ruefle, and Mark Zajicek. 2018. Organizational Models for Computer Security Incident Response Teams (CSIRTs). <https://doi.org/10.1184/R1/6575921.v1>

- [113] Iacovos Kirlappos, Simon Parkin, and M. Angela Sasse. 2014. Learning from “Shadow Security”: Why Understanding Non-compliance Provides the Basis for Effective Security. In *2014 IEEE Symposium on Security and Privacy Workshops (USEC 2014)*. IEEE, San Jose, CA, USA, n/a. <https://doi.org/10.14722/usec.2014.23007>
- [114] Iacovos Kirlappos, Simon Parkin, and Martina Angela Sasse. 2014. Learning from “Shadow Security”: Why Understanding Non-Compliance Provides the Basis for Effective Security. In *Workshop on Usable Security (USEC 2014)*. Internet Society, Reston, VA, USA, 1–8. <https://doi.org/10.14722/usec.2014.23007>
- [115] Laura Kocksch, Matthias Korn, Andreas Poller, and Susann Wagenknecht. 2018. Caring for IT security: Accountabilities, moralities, and oscillations in IT security practices. *Proceedings of the ACM on Human-Computer Interaction 2*, CSCW (2018), 1–20.
- [116] Angel Kodituwakku, Clark Xu, Daniel Rogers, David K. Ahn, and Errin W. Fulp. 2023. Temporal Aspects of Cyber Threat Intelligence. In *2023 IEEE International Conference on Big Data (BigData)*. IEEE, Los Alamitos, CA, USA, 6207–6211. <https://doi.org/10.1109/BigData59044.2023.10386664>
- [117] Faris Bugra Kokulu, Ananta Soneji, Tiffany Bao, Yan Shoshitaishvili, Ziming Zhao, Adam Doupé, and Gail-Joon Ahn. 2019. Matched and Mismatched SOCs: A Qualitative Study on Security Operations Center Issues. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. ACM, New York, NY, USA, 1955–1970.
- [118] Andrew Kompanek and Pawel Pawlinski. 2016. Evaluating Threat Intelligence Feeds FIRST Technical Colloquium for Threat Intelligence. <https://www.first.org/resources/papers/2016#munich2016>.
- [119] Arpine Korekryan. 2024. E-Government Survey 2024: Insights for ESCAP (AIS Steering Committee). [https://www.unescap.org/sites/default/d8files/event-documents/UN%20DESA%20E-gov%20Survey%202024%20Insights\\_for%20ESCAP%20event\\_24092024-ak.pdf](https://www.unescap.org/sites/default/d8files/event-documents/UN%20DESA%20E-gov%20Survey%202024%20Insights_for%20ESCAP%20event_24092024-ak.pdf) Accessed: 8-10-2025.
- [120] Platon Kotzias, Leyla Bilge, Pierre-Antoine Vervier, and Juan Caballero. 2019. Mind Your Own Business: A Longitudinal Study of Threats and Vulnerabilities in Enterprises. In *Proceedings 2019 Network and Distributed System Security Symposium*. Internet Society, San Diego, CA, USA, n/a. <https://doi.org/10.14722/ndss.2019.23522>
- [121] Sara Kraemer, Pascale Carayon, and John Clem. 2009. Human and organizational factors in computer and information security: Pathways to vulnerabilities. *Computers & Security* 28, 7 (2009), 509–520. <https://doi.org/10.1016/j.cose.2009.04.006>
- [122] Klaus Krippendorff. 2011. Computing Krippendorff’s Alpha-Reliability. <https://api.semanticscholar.org/CorpusID:59901023>

- [123] Katharina Krombholz, Wilfried Mayer, Martin Schmiedecker, and Edgar Weippl. 2017. "I have no idea what i'm doing": on the usability of deploying HTTPS. In *Proceedings of the 26th USENIX Conference on Security Symposium (SEC'17)*. USENIX Association, USA, 1339–1356.
- [124] Olaf Kruidhof. 2014. Evolution of National and Corporate CERTs – Trust, the Key Factor. In *Best Practices in Computer Network Defense: Incident Detection and Response*. IOS Press, Amsterdam, The Netherlands, 81–96. <https://doi.org/10.3233/978-1-61499-372-8-81>
- [125] Brenden Kuerbis and Farzaneh Badiei. 2017. Mapping the cybersecurity institutional landscape. *Digital Policy, Regulation and Governance* 19, 6 (2017), 466–492. <https://doi.org/10.1108/DPRG-05-2017-0024>
- [126] Marc Kühner, Christian Rossow, and Thorsten Holz. 2014. Paint It Black: Evaluating the Effectiveness of Malware Blacklists. In *Research in Attacks, Intrusions and Defenses*, Angelos Stavrou, Herbert Bos, and Georgios Portokalidis (Eds.). Springer International Publishing, Cham, 1–21.
- [127] Giuseppe Laurenza and Riccardo Lazzaretti. 2020. dAPTaset: A Comprehensive Mapping of APT-Related Data. In *Computer Security*, Apostolos P. Fournaris, Manos Athanatos, Konstantinos Lampropoulos, Sotiris Ioannidis, George Hatzivasilis, Ernesto Damiani, Habtamu Abie, Silvio Ranise, Luca Verderame, Alberto Siena, and Joaquin Garcia-Alfaro (Eds.). Springer International Publishing, Cham, 217–225.
- [128] Stevens Le Blond, Adina Uritesc, Cédric Gilbert, Zheng Leong Chua, Prateek Saxena, and Engin Kirda. 2014. A look at targeted attacks through the lense of an NGO. In *Proceedings of the 23rd USENIX Conference on Security Symposium (SEC'14)*. USENIX Association, USA, 543–558.
- [129] Beth L. Leech. 2002. Asking Questions: Techniques for Semistructured Interviews. *PS: Political Science & Politics* 35, 4 (Dec. 2002), 665–668. <https://doi.org/10.1017/S1049096502001129> Publisher: Cambridge University Press.
- [130] Alexander Leeuw. 2022. Geen extra geld voor uitvoering cybersecurity. <https://www.binnenlandsbestuur.nl/digitaal/geen-extra-financiële-middelen-voor-uitvoering-cybersecurity> Section: digitaal.
- [131] Cristoffer Leite, Jerry Den Hartog, and Daniel Ricardo dos Santos. 2024. Using DNS Patterns for Automated Cyber Threat Attribution. In *Proceedings of the 19th International Conference on Availability, Reliability and Security (ARES '24)*. Association for Computing Machinery, New York, NY, USA, Article 196, 11 pages. <https://doi.org/10.1145/3664476.3670870>
- [132] Rober Lemos. 2023. CISA Addresses 'Cyber Poor' Small Biz, Local Government. Online: <https://www.darkreading.com/threat-intelligence/cisa-addresses-cyber-poor-small-biz-local-government>. <https://www.darkreading.com/threat-intelligence/cisa-addresses-cyber-poor-small-biz-local-government>

- [//www.darkreading.com/threat-intelligence/cisa-addresses-cyber-poor-small-biz-local-government](http://www.darkreading.com/threat-intelligence/cisa-addresses-cyber-poor-small-biz-local-government) Accessed: 2023-05-16.
- [133] Frank Li, Zakir Durumeric, Jakub Czyz, Mohammad Karami, Michael Bailey, Damon McCoy, Stefan Savage, and Vern Paxson. 2016. You've Got Vulnerability: Exploring Effective Vulnerability Notifications. In *25th USENIX Security Symposium (USENIX Security 16)*. USENIX Association, Austin, TX, 1033–1050. <https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/li>
- [134] Frank Li, Grant Ho, Eric Kuan, Yuan Niu, Lucas Ballard, Kurt Thomas, Elie Bursztein, and Vern Paxson. 2016. Remediating Web Hijacking: Notification Effectiveness and Webmaster Comprehension. In *Proceedings of the 25th International World Wide Web Conference (WWW 2016)*. ACM, New York, NY, USA, 1007–1016.
- [135] Frank Li, Lisa Rogers, Arunesh Mathur, Nathan Malkin, and Marshini Chetty. 2019. Keepers of the Machines: Examining How System Administrators Manage Software Updates For Multiple Machines. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*. USENIX Association, Santa Clara, CA, 273–288. <https://www.usenix.org/conference/soups2019/presentation/li>
- [136] Vector Guo Li, Matthew Dunn, Paul Pearce, Damon McCoy, Geoffrey M. Voelker, and Stefan Savage. 2019. Reading the Tea leaves: A Comparative Analysis of Threat Intelligence. In *28th USENIX Security Symposium (USENIX Security 19)*. USENIX Association, Santa Clara, CA, 851–867. <https://www.usenix.org/conference/usenixsecurity19/presentation/li>
- [137] Michael Lipsky. 1980. *Street-Level Bureaucracy: Dilemmas of the Individual in Public Services*. Russell Sage Foundation, New York, NY, USA. <http://www.jstor.org/stable/10.7758/9781610447713>
- [138] Chuanying Lu and Luyao Zhang. 2022. A Chinese Perspective on Public Cyber Attribution. *China Quarterly of International Strategic Studies* 08, 01 (2022), 61–77. <https://doi.org/10.1142/S2377740022500026> arXiv:<https://doi.org/10.1142/S2377740022500026>
- [139] Gordon Lyon. 2023. Nmap: the Network Mapper - Free Security Scanner. <https://nmap.org/>
- [140] Minzhao Lyu, Hassan Habibi Gharakheili, and Vijay Sivaraman. 2022. Classifying and tracking enterprise assets via dual-grained network behavioral analysis. *Computer Networks* 218 (2022), 109387. <https://doi.org/10.1016/j.comnet.2022.109387>
- [141] Chanel Macabante, Sherry Wei, and David Schuster. 2019. Elements of Cyber-Cognitive Situation Awareness in Organizations. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 63, 1 (2019), 1624–1628. <https://doi.org/10.1177/1071181319631483> arXiv:<https://doi.org/10.1177/1071181319631483>

- [142] Stuart Madnick, Xitong Li, and Nazli Choucri. 2009. Experiences and Challenges with Using CERT Data to Analyze International Cyber Security. <https://doi.org/10.2139/ssrn.1478206>
- [143] Federico Maggi, Andrea Bellini, Guido Salvaneschi, and Stefano Zanero. 2011. Finding non-trivial malware naming inconsistencies. In *Proceedings of the 7th International Conference on Information Systems Security (ICISS'11)*. Springer-Verlag, Berlin, Heidelberg, 144–159. [https://doi.org/10.1007/978-3-642-25560-1\\_10](https://doi.org/10.1007/978-3-642-25560-1_10)
- [144] Christos Makridis, Lennart Maschmeyer, and Max Smeets. 2024. If it bleeps it leads? Media coverage on cyber conflict and misperception. *Journal of Peace Research* 61, 1 (2024), 72–86. <https://doi.org/10.1177/00223433231220264> arXiv:<https://doi.org/10.1177/00223433231220264>
- [145] Malpedia. 2025. Malpedia - Actors. [https://malpedia.caad.fkie.fraunhofer.de/actors?utm\\_source=chatgpt.comhttps://apt.etchda.or.th/cgi-bin/listgroups.cgi](https://malpedia.caad.fkie.fraunhofer.de/actors?utm_source=chatgpt.comhttps://apt.etchda.or.th/cgi-bin/listgroups.cgi) Accessed: 2025-08-06.
- [146] Mandiant. 2025. Google Threat Intelligence – know who’s targeting you. Google Cloud. <https://cloud.google.com/security/products/threat-intelligence> Accessed: 2025-07-07.
- [147] Shirang Mare, Mary Baker, and Jeremy Gummesson. 2016. A Study of Authentication in Daily Life. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. USENIX Association, Denver, CO, 189–206. <https://www.usenix.org/conference/soups2016/technical-sessions/presentation/mare>
- [148] Morgan Marquis-Boire, Marion Marschalek, and Claudio Guarnieri. 2015. Big Game Hunting: The Peculiarities in Nation-State Malware Research. *Black Hat 2015*, 1 (2015), 1–15. <https://www.blackhat.com/docs/us-15/materials/us-15-MarquisBoire-Big-Game-Hunting-The-Peculiarities-Of-Nation-State-Malware-Research.pdf> Black Hat USA, Las Vegas, NV, USA.
- [149] Giacomo Marzi, Marco Balzano, and Davide Marchiori. 2024. K-Alpha Calculator–Krippendorff’s Alpha Calculator: A user-friendly tool for computing Krippendorff’s Alpha inter-rater reliability coefficient. *MethodsX* 12 (2024), 102545. <https://doi.org/10.1016/j.mex.2023.102545>
- [150] John Matherly. 2015. *Complete Guide to Shodan*. Leanpub, Victoria, BC, Canada. <https://leanpub.next/shodan>
- [151] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. Reliability and Inter-rater Reliability in Qualitative Research: Norms and Guidelines for CSCW and HCI Practice. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (Nov. 2019), 1–23. <https://doi.org/10.1145/3359174>

- [152] Leigh Metcalf and Jonathan M. Spring. 2015. Blacklist Ecosystem Analysis: Spanning Jan 2012 to Jun 2014. In *Proceedings of the 2nd ACM Workshop on Information Sharing and Collaborative Security (WISCS '15)*. Association for Computing Machinery, New York, NY, USA, 13–22. <https://doi.org/10.1145/2808128.2808129>
- [153] Josh Meyer. 2022. Local governments are more vulnerable to cyberattacks than ever before. DHS wants mayors to step up. Online: <https://www.usatoday.com/story/news/politics/2022/02/08/local-government-cybersecurity-digital-threats/9208951002/>. <https://www.usatoday.com/story/news/politics/2022/02/08/local-government-cybersecurity-digital-threats/9208951002/> Accessed: 2023-05-24.
- [154] Adam Meyers. 2025. CrowdStrike and Microsoft Unite to Deconflict Cyber Threat Attribution. <https://www.crowdstrike.com/en-us/blog/crowdstrike-and-microsoft-unite-to-deconflict-cyber-threat-attribution/> Accessed: 2025-06-11.
- [155] Ola Aleksandra Michalec, Dirk van der Linden, Sveta Milyaeva, and Awais Rashid. 2020. Industry Responses to the European Directive on Security of Network and Information Systems (NIS): Understanding Policy Implementation Practices across Critical Infrastructures. In *Proceedings of the Sixteenth Symposium on Usable Privacy and Security (SOUPS 2020)*. USENIX Association, Berkeley, CA, USA, 301–317. <https://www.usenix.org/conference/soups2020/presentation/michalec>
- [156] Microsoft. 2025. How Microsoft names threat actors - Unified security operations. <https://learn.microsoft.com/en-us/unified-secops-platform/microsoft-threat-actor-naming>
- [157] Omid Mirzaei, Roman Vasilenko, Engin Kirda, Long Lu, and Amin Kharraz. 2021. SCRUTINIZER: Detecting Code Reuse in Malware via Decompilation and Machine Learning. In *Detection of Intrusions and Malware, and Vulnerability Assessment*, Leyla Bilge, Lorenzo Cavallaro, Giancarlo Pellegrino, and Nuno Neves (Eds.). Springer International Publishing, Cham, 130–150.
- [158] MISP. 2025. MISP Galaxy with Various Threat Intelligence Producers. <https://misp-galaxy.org/producer/> Accessed: 2025-08-07.
- [159] MITRE. 2025. MITRE ATT&CK®. <https://attack.mitre.org/>
- [160] MITRE. 2025. Tracking UNC2452-Related Reporting. <https://github.com/center-for-threat-informed-defense/public-resources/blob/master/solorigate/README.md> Accessed: 2025-08-12.
- [161] Aziz Mohaisen and Omar Alrawi. 2014. AV-Meter: An Evaluation of Antivirus Scans and Labels. In *Detection of Intrusions and Malware, and Vulnerability Assessment*, Sven Dietrich (Ed.). Springer International Publishing, Cham, 112–131.

- [162] S. R. B. Mohd Kassim, S. Li, and B. Arief. 2023. Understanding How National CSIRTs Evaluate Cyber Incident Response Tools and Data: Findings from Focus Group Discussions. *Digital Threats: Research and Practice* 4, 3 (2023), 18. <https://doi.org/10.1145/3609230>
- [163] Tyler Moore, Scott Dynes, and Frederick R. Chang. 2016. *Identifying How Firms Manage Cybersecurity Investment*. Technical Report. Workshop on the Economics of Information Security (WEIS). Available at <https://tylermoore.utulsa.edu/weis16ciso.pdf>.
- [164] Tyler W. Moore. 2008. *Cooperative attack and defense in distributed networks*. Technical Report 718. University of Cambridge, Computer Laboratory, Cambridge, United Kingdom. <https://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-718.pdf>
- [165] Stuart Murdoch and Nick Leaver. 2015. Anonymity vs. Trust in Cyber-Security Collaboration. In *Proceedings of the 2nd ACM Workshop on Information Sharing and Collaborative Security (WISCS '15)*. Association for Computing Machinery, New York, NY, USA, 27–29. <https://doi.org/10.1145/2808128.2808134>
- [166] Ali Naseri and Omid Azmoon. 2012. Proposition of model for CSIRT: Case study of telecommunication company in a province of Iran. *International Journal of Computer Science Issues (IJCSI)* 9, 1 (2012), 156.
- [167] Nationaal Cyber Security Centrum (NCSC-NL). n.d.. Over Mijncsc. <https://www.ncsc.nl/aansluiten-en-samenwerken/mijncsc/over-mijncsc>. Accessed September 11, 2025.
- [168] Nationaal Cybersecurity Centrum (NCSC). 2025. Ondersteuning bij cyberincidenten – Als sectoraal CSIRT. <https://www.ncsc.nl/documenten/factsheets/2025/februari/11/ondersteuning-bij-cyberincidenten---voor-nis2-organisaties>. Accessed April 22, 2025.
- [169] National Cyber Security Centre (NCSC). 2025. *Active Cyber Defence 2.0: Attack Surface Management Experiment Report*. Technical Report. National Cyber Security Centre, United Kingdom. <https://www.ncsc.gov.uk/files/Active-Cyber-Defence-2-ASM-experiment.pdf> Crown copyright 2025, accessed via <https://www.ncsc.gov.uk/files/Active-Cyber-Defence-2-ASM-experiment.pdf>.
- [170] NCSC.gov.uk. 2023. Early Warning - NCSC. Online: <https://www.earlywarning.service.ncsc.gov.uk/>. <https://www.earlywarning.service.ncsc.gov.uk/> Accessed: 2023-09-6.
- [171] Palo Alto Networks. 2025. Feed MISP Threat Actors – Cortex XSOAR Integration. [https://xsoar.pan.dev/docs/reference/integrations/feed-misp-threat-actors?utm\\_source=chatgpt.com](https://xsoar.pan.dev/docs/reference/integrations/feed-misp-threat-actors?utm_source=chatgpt.com). Accessed: 2025-11-11.
- [172] Katie Nickels. 2019. Cyber Indictments and Threat Intel: Why You Should Care. <https://medium.com/katies-five-cents/cyber-indictments-and-threat-intel-why-you-should-care-6336a14bb527>

- [173] NIST. 2013. Cybersecurity Framework. Online: <https://www.nist.gov/cyberframework>. <https://www.nist.gov/cyberframework> Accessed: 2023-06-2.
- [174] NIST. 2023. National Vulnerability Database - Home. Online: <https://nvd.nist.gov/>. <https://nvd.nist.gov/> Accessed: 2023-05-17.
- [175] NIST. 2023. National Vulnerability Database - Vulnerability Metrics. Online: <https://nvd.nist.gov/vuln-metrics/cvss>. <https://nvd.nist.gov/vuln-metrics/cvss> Accessed: 2023-05-17.
- [176] Donald Norris. 2021. A Look at Local Government Cybersecurity in 2020. Online: <https://icma.org/articles/pm-magazine/look-local-government-cybersecurity-2020>. <https://icma.org/articles/pm-magazine/look-local-government-cybersecurity-2020> Accessed: 2023-05-24.
- [177] Donald F Norris, Laura Mateczun, Anupam Joshi, and Tim Finin. 2021. Managing cybersecurity at the grassroots: Evidence from the first nationwide survey of local government cybersecurity. *Journal of Urban Affairs* 43, 8 (2021), 1173–1195.
- [178] Donald F. Norris, Laura K. Mateczun, and Richard F. Forno. 2022. *Cybersecurity and Local Government*. John Wiley & Sons, Hoboken, NJ, USA.
- [179] NOS Nieuws. 2021. Hack bij gemeente Hof van Twente veroorzaakt door te simpel wachtwoord. Online: <https://nos.nl/artikel/2372868-hack-bij-gemeente-hof-van-twente-veroorzaakt-door-te-simpel-wachtwoord>. <https://nos.nl/artikel/2372868-hack-bij-gemeente-hof-van-twente-veroorzaakt-door-te-simpel-wachtwoord> Accessed: 2023-05-11.
- [180] Justin Novak, Brittany Manley, David McIntire, Sharon Mudd, Angel Hueca, and Tracy Bills. 2021. The Sector CSIRT Framework: Developing Sector-Based Incident Response Capabilities. Carnegie Mellon University, Software Engineering Institute’s Digital Library. <https://doi.org/10.1184/R1/13624148> Accessed: 2024-Sep-23.
- [181] Justin Novak, Brittany Manley, David McIntire, Sharon Mudd, Angel Hueca, and Tracy Bills. 2021. The Sector CSIRT Framework: Developing Sector-Based Incident Response Capabilities.
- [182] Bank of England. 2024. CBEST Threat Intelligence-Led Assessments. <https://www.bankofengland.co.uk/financial-stability/operational-resilience-of-the-financial-sector/cbest-threat-intelligence-led-assessments-implementation-guide> Accessed: 2025-07-07.
- [183] National Institute of Standards and Technology. 2024. *The NIST Cybersecurity Framework (CSF) 2.0*. Technical Report CSWP.29. NIST. <https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.29.pdf> Accessed: 29-9-2025; DOI: 10.6028/NIST.CSWP.29.

- [184] National Institute of Standards and Technology. 2024. Perspectives on the CSF 1.1. <https://www.nist.gov/cyberframework/perspectives>. Accessed: 29-9-2025.
- [185] The White House Office of the Press Secretary. 2016. Presidential Policy Directive – United States Cyber Incident Coordination. <https://obamawhitehouse.archives.gov/the-press-office/2016/07/26/presidential-policy-directive-united-states-cyber-incident>
- [186] Jamie O’Hare, Rich Macfarlane, and Owen Lo. 2019. Identifying Vulnerabilities Using Internet-Wide Scanning Data. In *2019 IEEE 12th International Conference on Global Security, Safety and Sustainability (ICGS3)*. IEEE, Porto, Portugal, 1–10. <https://doi.org/10.1109/ICGS3.2019.8688018>
- [187] Council on Foreign Relations. 2022. Targeting of Ukrainian local government organization. <https://www.cfr.org/cyber-operations/targeting-ukrainian-local-government-organization>. Accessed: 1-10-2025.
- [188] Talha Ongun, Jack W. Stokes, Jonathan Bar Or, Ke Tian, Farid Tajaddodianfar, Joshua Neil, Christian Seifert, Alina Oprea, and John C. Platt. 2021. Living-Off-The-Land Command Detection Using Active Learning. In *International Symposium on Research in Attacks, Intrusions and Defenses (RAID)*. Springer, Cham, Switzerland, 225–245.
- [189] Anthony J. Onwuegbuzie, Wendy B. Dickinson, Nancy L. Leech, and Annmarie G. Zoran. 2009. A Qualitative Framework for Collecting and Analyzing Data in Focus Group Research. *International Journal of Qualitative Methods* 8, 3 (2009), 1–21. <https://doi.org/10.1177/160940690900800301> arXiv:<https://doi.org/10.1177/160940690900800301>
- [190] Andreas Oster, Eivor Wiking, Gunnar H. Nilsson, and Christina B. Olsson. 2024. Patients’ expectations of primary health care from both patients’ and physicians’ perspectives: a questionnaire study with a qualitative approach. *BMC Primary Care* 25, 1 (April 2024), 128. <https://doi.org/10.1186/s12875-024-02389-2>
- [191] Ruoming Pang, Mark Allman, Mike Bennett, Jason Lee, Vern Paxson, and Brian Tierney. 2005. A First Look at Modern Enterprise Traffic. In *Proceedings of the 5th ACM SIGCOMM Conference on Internet Measurement (IMC 2005)*. USENIX Association, Berkeley, CA, USA, 2–13. <https://www.usenix.org/conference/imc-05/first-look-modern-enterprise-traffic>
- [192] Eric Pauley, Paul Barford, and Patrick McDaniel. 2023. The CVE Wayback Machine: Measuring Coordinated Disclosure from Exploits against Two Years of Zero-Days. In *Proceedings of the 2023 ACM on Internet Measurement Conference (IMC ’23)*. Association for Computing Machinery, New York, NY, USA, 236–252. <https://doi.org/10.1145/3618257.3624810>

- [193] Martin Plattner and Rainer Böhme. 2023. *More Security, Less Harm? Exploring the Link between Security Measures and Direct Costs of Cyber Incidents within Firms Using PLS-PM*. Working Paper. University of Tübingen. <https://bibliographie.uni-tuebingen.de/xmlui/bitstream/handle/10900/142498/More%20Security.pdf> Accessed: 2025-01-31.
- [194] Kevin Poireault. 2023. Understanding Threat Actor Naming Conventions. <https://www.infosecurityeurope.com/en-gb/blog/threat-vectors/understanding-threat-actor-naming-conventions.html> Accessed: 2025-07-30.
- [195] Walter W. Powell. 1990. Neither Market Nor Hierarchy: Network Forms of Organization. In *Research in Organizational Behavior*, Barry M. Staw and L. L. Cummings (Eds.). Vol. 12. JAI Press, Greenwich, CT, 295–336.
- [196] The Canadian Press . 2018. Ontario police warn of recent cyberattacks targeting local governments | CBC News. Online: <https://www.cbc.ca/news/canada/toronto/cyberattacks-targeting-local-government-ontario-1.4824772>. Accessed: 2023-05-11.
- [197] PriceWaterhouseCoopers. 2024. How to mature your organisation's threat intelligence capabilities. <https://www.pwc.com/gx/en/issues/cybersecurity/cyber-threat-intelligence/maturing-your-threat-intelligence-capabilities.html>
- [198] Qualys. 2023. Qualys VMDR - Vulnerability Management Tool | Qualys. <https://www.qualys.com/apps/vulnerability-management-detection-response/> Accessed: 2023-09-20.
- [199] Sveriges Radio. 2021. Kalix municipality still recovering from cyber-attack. <https://www.sverigesradio.se/artikel/kalix-municipality-still-recovering-from-cyber-attack>. Accessed: 29-9-2025.
- [200] Prashanth Rajivan and Nancy Cooke. 2017. *Impact of team collaboration on cybersecurity situational awareness*. Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), Vol. 10030. Springer Verlag, Germany, 203–226. [https://doi.org/10.1007/978-3-319-61152-5\\_8](https://doi.org/10.1007/978-3-319-61152-5_8)
- [201] Rapid7. 2023. Metasploit. <https://rapid7.github.io/metasploit-framework/> Accessed: 2023-09-20.
- [202] RedDrip7. 2025. RedDrip7/APT\_Digital\_Weapon. [https://github.com/RedDrip7/APT\\_Digital\\_Weapon](https://github.com/RedDrip7/APT_Digital_Weapon) original-date: 2019-12-05T04:11:17Z.
- [203] Yitong Ren, Yanjun Xiao, Yinghai Zhou, Zhiyong Zhang, and Zhihong Tian. 2023. CSKG4APT: A Cybersecurity Knowledge Graph for Advanced Persistent Threat Organization Attribution. *IEEE Transactions on Knowledge and Data Engineering* 35, 6 (2023), 5695–5709. <https://doi.org/10.1109/TKDE.2022.3175719>

- [204] Christopher Rentrop and S. Zimmermann. 2012. Shadow IT – Management and Control of Unofficial IT. In *Proceedings of the 6th International Conference on Digital Society (ICDS 2012)*. IARIA, Valencia, Spain, 230–235. <https://www.semanticscholar.org/paper/Shadow-IT-Management-and-Control-of-Unofficial-IT-Rentrop-Zimmermann/609ad294cdefbb66b7d97271ebc229eeab79c315>
- [205] Christopher Rentrop and Stephan Zimmermann. 2012. Shadow IT: Management and Control of Unofficial IT. In *Proceedings of the 6th International Conference on Digital Society (ICDS 2012)*. IARIA, Wilmington, DE, USA, 98–102.
- [206] Kristin Repchick, Stephen Zaccaro, Lois Tetrick, Julie Steinke, Daniel Shore, Carolyn Winslow, Amber reecho, Hargrove, Balca Alaybek, Jennifer Green, Tracy McCausland, and Alan Tomassetti. 2016. Improving social maturity of cybersecurity incident response teams.
- [207] Reuters. 2018. Atlanta ransomware attack throws city services into disarray. <https://www.reuters.com/article/usa-georgia-cyber/atlanta-ransomware-attack-throws-city-services-into-disarray-idUSL1N1R51V9/>. Accessed: 29-9-2025.
- [208] Reuters. 2019. Johannesburg power body hit by ransomware attack. <https://www.reuters.com/article/world/johannesburg-power-body-hit-by-ransomware-attack-idUSKCN1UK15G/>. Accessed: YYYY-MM-DD.
- [209] Thomas Rid and Ben Buchanan. 2015. Attributing Cyber Attacks. *Journal of Strategic Studies* 38, 1-2 (Jan. 2015), 4–37. <https://doi.org/10.1080/01402390.2014.977382> Publisher: Routledge \_eprint: <https://doi.org/10.1080/01402390.2014.977382>.
- [210] Thea Riebe, Marc-André Kaufhold, and Christian Reuter. 2021. The Impact of Organizational Structure and Technology Use on Collaborative Practices in Computer Emergency Response Teams: An Empirical Study. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 478 (Oct. 2021), 30 pages. <https://doi.org/10.1145/3479865>
- [211] Sasha Romanosky and Benjamin Boudreaux. 2021. Private-Sector Attribution of Cyber Incidents: Benefits and Risks to the U.S. Government. *International Journal of Intelligence and CounterIntelligence* 34, 3 (2021), 463–493. <https://doi.org/10.1080/08850607.2020.1783877> arXiv:<https://doi.org/10.1080/08850607.2020.1783877>
- [212] Ishai Rosenberg, Guillaume Sicard, and Eli (Omid) David. 2017. DeepAPT: Nation-State APT Attribution Using End-to-End Deep Neural Networks. In *Artificial Neural Networks and Machine Learning – ICANN 2017*, Alessandra Lintas, Stefano Rovetta, Paul F.M.J. Verschure, and Alessandro E.P. Villa (Eds.). Springer International Publishing, Cham, 91–99.

- [213] Florian Roth. 2018. The Newcomer's Guide to Cyber Threat Actor Naming. <https://cyb3rops.medium.com/the-newcomers-guide-to-cyber-threat-actor-naming-7428e18ee263> Accessed: 2025-06-10.
- [214] Vinay Sachidananda, Rajendra Patil, Akshay Sachdeva, Kwok-Yan Lam, and Liu Yang. 2023. APTer: Towards the Investigation of APT Attribution. In *2023 IEEE Conference on Dependable and Secure Computing (DSC)*. IEEE, Los Alamitos, CA, USA, 1–10. <https://doi.org/10.1109/DSC61021.2023.10354155>
- [215] Aakanksha Saha, Jorge Blasco, Lorenzo Cavallaro, and Martina Lindorfer. 2024. ADAPT it! Automating APT Campaign and Group Attribution by Leveraging and Linking Heterogeneous Files. In *Proceedings of the 27th International Symposium on Research in Attacks, Intrusions and Defenses (RAID '24)*. Association for Computing Machinery, New York, NY, USA, 114–129. <https://doi.org/10.1145/3678890.3678909>
- [216] Aakanksha Saha, Lorenzo Cavallaro, James Mattei, Daniel Votipka, Jorge Blasco, and Martina Lindorfer. 2025. Expert Insights into Advanced Persistent Threats: Analysis, Attribution, and Challenges.
- [217] Aakanksha Saha, Martina Lindorfer, and Juan Caballero. 2025. From IOCs to Group Profiles: On the Specificity of Threat Group Behaviors in CTI Knowledge Bases. arXiv:cs.CR/2506.10645 <https://arxiv.org/abs/2506.10645>
- [218] Emmanouil Samanis, Joseph Gardiner, and Awais Rashid. 2022. SoK: A Taxonomy for Contrasting Industrial Control Systems Asset Discovery Tools. In *Proceedings of the 17th International Conference on Availability, Reliability and Security (ARES '22)*. Association for Computing Machinery, New York, NY, USA, Article 28, 12 pages. <https://doi.org/10.1145/3538969.3538979>
- [219] Mario Saraiva and Nuno Mateus-Coelho. 2022. CyberSoc Framework a Systematic Review of the State-of-Art. *Procedia Computer Science* 204 (2022), 961–972. <https://doi.org/10.1016/j.procs.2022.08.117> Publisher Copyright: © 2022 Elsevier B.V.. All rights reserved.; 2022 International Conference on Industry Sciences and Computer Science Innovation, iSCSi 2022 ; Conference date: 09-03-2022 Through 11-03-2022.
- [220] Thomas Schaberreiter, Veronika Kupfersberger, Konstantinos Rantos, Arnolnt Spyros, Alexandros Papanikolaou, Christos Ilioudis, and Gerald Quirchmayr. 2019. A Quantitative Evaluation of Trust in the Quality of Cyber Threat Intelligence Sources. In *Proceedings of the 14th International Conference on Availability, Reliability and Security (ARES '19)*. Association for Computing Machinery, New York, NY, USA, Article 83, 10 pages. <https://doi.org/10.1145/3339252.3342112>
- [221] Daniel Schlette, Fabian Böhm, Marco Caselli, and Günther Pernul. 2021. Measuring and visualizing cyber threat intelligence quality. *International Journal of Information Security* 20, 1 (Feb. 2021), 21–38. <https://doi.org/10.1007/s10207-020-00490-y>

- [222] Maurice Schubert and Yasser Aboukir. 2025. Threat-Led Penetration Testing: A proactive approach to cybersecurity | Deloitte Luxembourg | Future of Advice. <https://www.deloitte.com/lu/en/our-thinking/future-of-advice/threat-led-penetration-testing-proactive-approach-cybersecurity.html> accessed 2025-07-07.
- [223] Kaspersky Securelist. 2025. Targeted cyberattacks logbook. <https://apt.securelist.com>
- [224] Multi-State Information Sharing and Analysis Center (MS-ISAC). 2015. What is Cyber Threat Intelligence? <https://www.cisecurity.org/blog/what-is-cyber-threat-intelligence/>
- [225] Erika Stanish. 2023. Municipal Water Authority of Aliquippa hacked by Iranian-backed cyber group. <https://www.cbsnews.com/pittsburgh/news/municipal-water-authority-of-aliquippa-hacked-iranian-backed-cyber-group/>. Accessed: 1-10-2025.
- [226] T. Steffens. 2020. *Attribution of Advanced Persistent Threats: How to Identify the Actors Behind Cyber-Espionage*. Springer Berlin Heidelberg, Berlin, Germany. <https://books.google.nl/books?id=6FryDwAAQBAJ>
- [227] Vilja Steffensen and Vahiny Gnanasekaran. 2024. Information Sharing between the Computer Security Incident Response Team and its Members: An Empirical Study. *Norsk IKT-konferanse for forskning og utdanning* 3, 3 (Nov. 2024), n/a. <https://www.ntnu.no/ojs/index.php/nikt/article/view/6250>
- [228] Julie Steinke, Balca Bolunmez, Laura Fletcher, Vicki Wang, Alan J. Tomassetti, Kristin M. Repchick, Stephen J. Zaccaro, Reeshad S. Dalal, and Lois E. Tetrick. 2015. Improving Cybersecurity Incident Response Team Effectiveness Using Teams-Based Research. *IEEE Security and Privacy* 13, 4 (jul 2015), 20–29. <https://doi.org/10.1109/MSP.2015.71>
- [229] Brent Stephens, Alan L. Cox, Scott Rixner, and T. S. Eugene Ng. 2011. A Scalability Study of Enterprise Network Architectures. In *Proceedings of the 2011 ACM/IEEE Seventh Symposium on Architectures for Networking and Communications Systems (ANCS '11)*. IEEE Computer Society, USA, 111–121. <https://doi.org/10.1109/ANCS.2011.28>
- [230] Don Stikvoort. 2015. SIM3: Security Incident Management Maturity Model. <https://cybilportal.org/publications/sim3-security-incident-management-maturity-model/> Accessed: 2024-09-23.
- [231] Ben Stock, Giancarlo Pellegrino, Christian Rossow, Martin Johns, and Michael Backes. 2016. Hey, you have a problem: on the feasibility of large-scale web vulnerability notification. In *Proceedings of the 25th USENIX Conference on Security Symposium (SEC'16)*. USENIX Association, USA, 1015–1032.

- [232] Sathya Chandran Sundaramurthy, John McHugh, Xinming Simon Ou, S. Raj Rajagopalan, and Michael Wesch. 2014. An Anthropological Approach to Studying CSIRTs. *IEEE Security & Privacy* 12, 05 (Sept. 2014), 52–60. <https://doi.org/10.1109/MSP.2014.84>
- [233] Deniz Susar. 2024. Local E-Government Development. In *E-Government Survey 2024: Accelerating Digital Transformation for Sustainable Development*. United Nations, New York, NY, USA, 135–156. <https://desapublications.un.org/sites/default/files/publications/2024-10/Chapter%204%20E-Government%20Survey%202024.pdf> Accessed: 8-10-2025.
- [234] Samaneh Tajalizadehkhoob. 2018. *The Role of Hosting Providers in Web Security: Understanding and Improving Security Incentives and Performance via Analysis of Large-scale Incident Data*. PhD dissertation. Delft University of Technology, Delft, The Netherlands. <https://doi.org/10.4233/uuid:c343a2dd-15d1-4921-9b45-f00ee38177d8>
- [235] Samaneh Tajalizadehkhoob, Tom Van Goethem, Maciej Korczyński, Arman Noroozian, Rainer Böhme, Tyler Moore, Wouter Joosen, and Michel van Eeten. 2017. Herding Vulnerable Cats: A Statistical Approach to Disentangle Joint Responsibility for Web Security in Shared Hosting. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS '17)*. Association for Computing Machinery, New York, NY, USA, 225–238. <https://doi.org/10.1145/3133956.3133971>
- [236] Godfrey Tan, Massimiliano Poletto, John Guttag, and Frans Kaashoek. 2003. Role Classification of Hosts Within Enterprise Networks. In *2003 USENIX Annual Technical Conference (USENIX ATC '03)*. USENIX Association, San Antonio, TX, 15–28. [https://www.usenix.org/legacy/event/usenix03/tech/full\\_papers/full\\_papers/tan/tan.pdf](https://www.usenix.org/legacy/event/usenix03/tech/full_papers/full_papers/tan/tan.pdf)
- [237] Abbas Tashakkori and Charles Teddlie. 2003. *Handbook of Mixed Methods in Social & Behavioral Research*. SAGE Publications, Thousand Oaks, CA, USA. <https://books.google.nl/books?id=F8BF0M8DCKoC>
- [238] Chris Teale. 2023. Southern states have the most open cyber exposures, report finds. Online: <https://gcn.com/cybersecurity/2023/02/southern-states-have-most-open-cyber-exposures-report-finds/383418/>. <https://gcn.com/cybersecurity/2023/02/southern-states-have-most-open-cyber-exposures-report-finds/383418/> Accessed: 2023-05-24.
- [239] Tenable. 2023. CVSS Scores vs. VPR (Nessus 10.5). Online: <https://docs.tenable.com/nessus/Content/RiskMetrics.htm>. <https://docs.tenable.com/nessus/Content/RiskMetrics.htm> Accessed: 2023-05-2.
- [240] Tenable. 2023. Nessus Vulnerability Scanner. <https://www.tenable.com/lp/campaigns/22/nessus-multiprdct/buy/> Accessed: 2023-09-20.

- [241] Tim Tenbenschel. 2005. Multiple Modes of Governance: Disentangling the Alternatives to Hierarchies and Markets. *Public Management Review* 7, 2 (2005), 267–288.
- [242] The Shadowserver Foundation. n.d.. The Shadowserver Foundation. <https://www.shadowserver.org/>. Accessed April 28, 2025.
- [243] Kurt Thomas, Rony Amira, Adi Ben-Yoash, Ori Folger, Amir Hardon, Ari Berger, Elie Bursztein, and Michael Bailey (Eds.). 2016. *The Abuse Sharing Economy: Understanding the Limits of Threat Exchanges*.
- [244] Susan M. Tisdale. 2015. Cybersecurity: Challenges From a Systems, Complexity, Knowledge Management and Business Intelligence Perspective. In *Issues in Information Systems, Volume 16, Issue III*. International Academy of Computer Information Systems, Athens, GA, USA, 191–198. [https://iacis.org/iis/2015/3\\_iis\\_2015\\_191-198.pdf](https://iacis.org/iis/2015/3_iis_2015_191-198.pdf)
- [245] Giorgio Di Tizio, Michele Armellini, and Fabio Massacci. 2023. Software Updates Strategies: A Quantitative Evaluation Against Advanced Persistent Threats. *IEEE Transactions on Software Engineering* 49, 3 (2023), 1359–1373. <https://doi.org/10.1109/TSE.2022.3176674>
- [246] Breno Tostes, Leonardo Ventura, Enrico Lovat, Matheus Martins, and Daniel Menasché. 2023. Learning When to Say Goodbye: What Should Be the Shelf Life of an Indicator of Compromise?. In *2023 IEEE International Conference on Cyber Security and Resilience (CSR)*. IEEE, Los Alamitos, CA, USA, 503–510. <https://doi.org/10.1109/CSR57506.2023.10224937>
- [247] Home Office UK. 2019. Cyber Threat Intelligence in Government: A Guide for Decision Makers and Analysts. <https://hodigital.blog.gov.uk/wp-content/uploads/sites/161/2020/03/Cyber-Threat-Intelligence-A-Guide-For-Decision-Makers-and-Analysts-v2.0.pdf> Accessed: 2025-07-07.
- [248] NCSC UK. 2025. Find an assured Cyber Incident Response provider. <https://www.ncsc.gov.uk/schemes/cyber-incident-response/find-a-provider> Accessed: 2025-08-12.
- [249] European Union. 2025. Digital Operational Resilience Act (DORA) - EIOPA. [https://www.eiopa.europa.eu/digital-operational-resilience-act-dora\\_en](https://www.eiopa.europa.eu/digital-operational-resilience-act-dora_en) Accessed: 2025-07-07.
- [250] Secureworks Counter Threat Unit. 2023. Iranian Cyber Av3ngers Compromise Unitronics Systems. <https://www.secureworks.com/blog/iranian-cyber-av3ngers-compromise-unitronics-systems>. Accessed: 1-10-2025.
- [251] U.S. Department of Homeland Security. 2024. *FY2024 Annual Performance Report, Appendix D: Measure Descriptions, Data Collection Methodologies, and Completeness and Reliability Information*. Annual Performance Report. U.S. Department of Homeland Security. [https://www.dhs.gov/sites/default/files/2025-01/2025\\_0117\\_dhs\\_annual\\_performance\\_report\\_fy2024\\_appendixd.pdf](https://www.dhs.gov/sites/default/files/2025-01/2025_0117_dhs_annual_performance_report_fy2024_appendixd.pdf) Appendix D, p. 25.

- [252] Christine Utz, Matthias Michels, Martin Degeling, Ninja Marnau, and Ben Stock. 2023. Comparing Large-Scale Privacy and Security Notifications. *Proceedings on Privacy Enhancing Technologies* 2023, 3 (July 2023), 173–193. <https://publications.cispa.saarland/3918/> ISSN: 2299-0984.
- [253] Steven Van de Walle and Geert Bouckaert. 2003. *Public Service Performance and Trust in Government: The Problem of Causality*. Eburon Academic Publishers, Delft, The Netherlands. [https://repub.eur.nl/pub/41526/Metis\\_173779.pdf](https://repub.eur.nl/pub/41526/Metis_173779.pdf) Accessed: 2025-11-09.
- [254] Max van der Horst, Ricky Kho, Olga Gadyatskaya, Michel Mollema, Michel Van Eeten, and Yury Zhauniarovich. 2025. High Stakes, Low Certainty: Evaluating the Efficacy of High-Level Indicators of Compromise in Ransomware Attribution | USENIX. <https://www.usenix.org/conference/usenixsecurity25/presentation/van-der-horst>
- [255] Rick van der Kleij, Geert Kleinhuis, and Heather Young. 2017. Computer Security Incident Response Team Effectiveness: A Needs Assessment. *Frontiers in Psychology* 8 (2017), 194. <https://www.frontiersin.org/articles/10.3389/fpsyg.2017.00194>
- [256] Veerle van Harten, Carlos Hernandez Ganan, Michel van Eeten, and Simon Parkin. 2025. “All Sorts of Other Reasons to Do It”: Explaining the Persistence of Sub-optimal IoT Security Advice. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 387, 19 pages. <https://doi.org/10.1145/3706598.3713719>
- [257] Koen van Hove, Jeroen van der Ham-de Vos, and Roland van Rijswijk-Deij. 2023. Your Vulnerability Disclosure Is Important To Us: An Analysis of Coordinated Vulnerability Disclosure Responses Using a Real Security Issue. arXiv:cs.NI/2312.07284 <https://arxiv.org/abs/2312.07284>
- [258] Vereniging van Nederlandse Gemeenten (VNG). 2021. *Bestuurlijke lessen uit de hack bij Hof van Twente*. Technical Report. VNG. <https://vng.nl/sites/default/files/2021-04/20210401-bestuurlijke-lessen-uit-de-hack-bij-hof-van-twente.pdf> PDF; accessed via <https://vng.nl/sites/default/files/2021-04/20210401-bestuurlijke-lessen-uit-de-hack-bij-hof-van-twente.pdf>.
- [259] Kelli Vanderlee. 2020. DebUNCing Attribution: How Mandiant Tracks Uncategorized Threat Actors | Mandiant. <https://cloud.google.com/blog/topics/threat-intelligence/how-mandiant-tracks-uncategorized-threat-actors> Accessed: 2025-06-12.
- [260] Nicole F. Velasquez and Suzanne P. Weisband. 2009. System administrators as broker technicians. In *Proceedings of the Symposium on Computer Human Interaction for the Management of Information Technology (CHI/MIT '09)*. Association

- for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/1641587.1641588>
- [261] Mathew Vermeer, Natalia Kadenko, Michel van Eeten, Carlos Gañán, and Simon Parkin. 2023. Alert Alchemy: SOC Workflows and Decisions in the Management of NIDS Rules. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (CCS '23)*. Association for Computing Machinery, New York, NY, USA, 2770–2784. <https://doi.org/10.1145/3576915.3616581>
- [262] Mathew Vermeer, Jonathan West, Alejandro Cuevas, Shuonan Niu, Nicolas Christin, Michel Van Eeten, Tobias Fiebig, Carlos Ganan, and Tyler Moore. 2021. SoK: A Framework for Asset Discovery: Systematizing Advances in Network Measurements for Protecting Organizations. In *2021 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, Vienna, Austria, 440–456. <https://doi.org/10.1109/EuroSP51992.2021.00037>
- [263] VirusTotal. 2025. VirusTotal — Upload and analyse files and URLs for free. <https://www.virustotal.com/gui/home/upload>. Accessed: 2025-11-11.
- [264] Daniel Votipka, Rock Stevens, Elissa Redmiles, Jeremy Hu, and Michelle Mazurek. 2018. Hackers vs. Testers: A Comparison of Software Vulnerability Discovery Processes. In *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, Piscataway, NJ, USA, 374–391. <https://doi.org/10.1109/SP.2018.00003> ISSN: 2375-1207.
- [265] Cynthia Wagner, Alexandre Dulaunoy, Gérard Wagener, and Andras Iklody. 2016. MISP: The Design and Implementation of a Collaborative Threat Intelligence Sharing Platform. In *Proceedings of the 2016 ACM Workshop on Information Sharing and Collaborative Security*. ACM, New York, NY, USA, 49–56.
- [266] Qinqin Wang, Hanbing Yan, and Zhihui Han. 2021. Explainable APT Attribution for Malware Using NLP Techniques. In *2021 IEEE 21st International Conference on Software Quality, Reliability and Security (QRS)*. IEEE, Los Alamitos, CA, USA, 70–80. <https://doi.org/10.1109/QRS54544.2021.00018>
- [267] YM Wara and D Singh. 2015. A guide to establishing computer security incident response team (CSIRT) for national research and education network (NREN). *African Journal of Computing & ICT* 8, 2 (2015), 1–8.
- [268] Arun Warikoo. 2021. The Triangle Model for Cyber Threat Attribution. *Journal of Cyber Security Technology* 5 (03 2021), 1–18. <https://doi.org/10.1080/23742917.2021.1895532>
- [269] Claudia Werker and Ward Ooms. 2020. Substituting face-to-face contacts in academics' collaborations: modern communication tools, proximity, and brokerage. *Studies in Higher Education* 45, 7 (2020), 1431–1447. <https://doi.org/10.1080/03075079.2019.1655723> arXiv:<https://doi.org/10.1080/03075079.2019.1655723>
- [270] Jonathan Codi West and Tyler Moore. 2022. Longitudinal Study of Internet-Facing OpenSSH Update Patterns. In *Passive and Active Measurement (Lecture Notes in*

- Computer Science*), Oliver Hohlfeld, Giovane Moura, and Cristel Pelsser (Eds.). Springer International Publishing, Cham, 675–689. [https://doi.org/10.1007/978-3-030-98785-5\\_30](https://doi.org/10.1007/978-3-030-98785-5_30)
- [271] Moira J. West-Brown, Don Stikvoort, Klaus-Peter Kossakowski, Georgia Killcrece, Robin Ruefle, and Mark Zajicek. 1998. *Handbook for Computer Security Incident Response Teams (CSIRTs)*. Carnegie Mellon University, Software Engineering Institute, Pittsburgh, PA, USA.
- [272] Johannes Wiik, Jose Gonzalez, and Klaus-Peter Kossakowski. 2006. Effectiveness of Proactive CSIRT Services. In *Proceedings of the TF-CSIRT Meeting 2006*. TERENA, Innsbruck, Austria, 67–81.
- [273] Chris Wilson. 2019. New Orleans Cyber Attack: When the City’s Sick of Ransomware. <https://time.com/5750242/new-orleans-cyber-attack/>. Accessed: 29-9-2025.
- [274] Daniel W. Woods and Rainer Böhme. 2021. Systematization of Knowledge: Quantifying Cyber Risk. In *Proceedings of the IEEE Symposium on Security and Privacy (SP)*. IEEE, San Francisco, CA, USA, 14. <https://doi.org/10.1109/SP40001.2021.00060>
- [275] Oliver Wright. 2020. Redcar and Cleveland council hit by cyber attack. <https://www.theguardian.com/technology/2020/feb/27/redcar-and-cleveland-council-hit-by-cyber-attack>. Accessed: 29-9-2025.
- [276] Julia Wunder, Alan Corona, Andreas Hammer, and Zinaida Benenson. 2024. On NVD Users’ Attitudes, Experiences, Hopes, and Hurdles. *Digital Threats* 5, 3, Article 33 (Oct. 2024), 19 pages. <https://doi.org/10.1145/3688806>
- [277] Nan Xiao, Bo Lang, Ting Wang, and Yikai Chen. 2024. APT-MMF: An advanced persistent threat actor attribution method based on multimodal and multilevel feature fusion. arXiv:cs.CR/2402.12743 <https://arxiv.org/abs/2402.12743>
- [278] Rei Yamagishi, Shota Fujii, Shingo Yasuda, Takayuki Sato, and Ayako A. Hasegawa. 2025. Collaborative Work in Malware Analysis: Understanding the Roles and Challenges of Malware Analysts. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI ’25)*. Association for Computing Machinery, New York, NY, USA, Article 865, 15 pages. <https://doi.org/10.1145/3706598.3713652>
- [279] Miuyin Yong Wong, Matthew Landen, Manos Antonakakis, Douglas M. Blough, Elissa M. Redmiles, and Mustaque Ahamad. 2021. An Inside Look into the Practice of Malware Analysis. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security (CCS ’21)*. Association for Computing Machinery, New York, NY, USA, 3053–3069. <https://doi.org/10.1145/3460120.3484759>

- [280] Orçun Çetin, Mohammad Hanif Jhaveri, Carlos Gañán, Michel van Eeten, and Tyler Moore. 2016. Understanding the role of sender reputation in abuse reporting and cleanup. *Journal of Cybersecurity* 2, 1 (12 2016), 83–98. <https://doi.org/10.1093/cybsec/tyw005> arXiv:<https://academic.oup.com/cybersecurity/article-pdf/2/1/83/10833175/tyw005.pdf>





# APPENDIX FOR CHAPTER 2

## A.1. RESPONDENT DETAILS

## A.2. INTERVIEW PROTOCOL

### INTRODUCTION AND BACKGROUND

- Can you tell me about yourself?
- What does a typical day look like?
- How many devices and servers are you managing, and how big is the team?
- How is security organized in your organization?
- What do you consider the biggest obstacles in security?
- What are the Internet-facing systems of the municipality?
- How do you monitor those systems?
- How do you stay up to date on vulnerabilities? Is that an active process?

### ADVISORIES AND CERT

- Who receives advisories and notifications from the CERT?
- Who manages the IP ranges, an individual or a team?
- Who is responsible for following up after a notification from the CERT?
- How do you determine if a notification is relevant?
- Do you report changes in your infrastructure to the CERT?

### SPECIFIC VULNERABLE SYSTEMS

- For system X, we detected service Y and version Z. Is that correct? Did you run a backport? Is it vulnerable? If so, which CVEs? How did you obtain that information?
- Are you aware of the latest version of service X? How do you obtain that information?
- How do you deal with those CVEs?
- Did you apply mitigation strategies? Why?
- Do you have other mitigation strategies that were not used? Do you have examples?
- How does the location or function of the system influence the choice of a mitigation strategy?

**Table A.1: Overview of the respondents per municipality**

Respondent Id	Muni Id	Role	Gender	Vulnerable Hosts
1	1	System Administrator	Male	3
2	2	System Administrator	Male	2
3	3	System Administrator	Male	1
4	3	Security Officer/Engineer	Male	1
5	4	(C)ISO	Male	1
6	5	System Administrator	Male	4
7	6	System Administrator	Male	5
8	6	(C)ISO	Male	5
9	7	Network Administrator	Male	5
10	7	(C)ISO	Male	5
11	7	Security Officer/Engineer	Male	5
12	8	Security Officer/Engineer	Male	7
13	9	Security Officer/Engineer	Male	2
14	10	Security Officer/Engineer	Male	0
15	10	Security Officer/Engineer	Male	0
16	11	Security Officer/Engineer	Male	0
17	11	System Administrator	Male	0
18	12	System Administrator	Male	1
19	12	System Administrator	Male	1
20	12	System Administrator	Male	1
21	13	Security Officer/Engineer	Male	0
22	13	(C)ISO	Female	0
23	13	(C)ISO	Male	0
24	14	Network Administrator	Male	1
25	14	Security Officer/Engineer	Male	1
26	15	(C)ISO	Male	0
27	16	Security Officer/Engineer	Female	0
28	16	Security Officer/Engineer	Male	0
29	16	Security Officer/Engineer	Male	0

## A.3. CODES

### DETERMINING VULNERABLE SYSTEMS

*Subcodes* Active Search; Asset Inventory; Notifications; Not determined; Vulnerability Scanning.

### MITIGATION STRATEGIES

*Subcodes* Strategies Notifier; Security Tools; Isolating systems; Managed Services; Non-internet facing.

### PRIORITIZATION AND PATCH PRACTICES

*Subcodes* Critical versus non-critical score; Internet-facing vs non-Internet-facing; Latest version on install.

### RESPONSIBILITIES

*Subcodes* Compliance; Dependencies in products; External, Cloud and SaaS services; Public IPs and internal network; Security role in organization; Users and awareness.

### UNDERMINING SECURITY

*Subcodes* Capacity for security tasks; Budget; People and skills; Priorities, partners and collaborations; Legacy systems.

## A.4. VENN DIAGRAMS TOTAL IP SETS

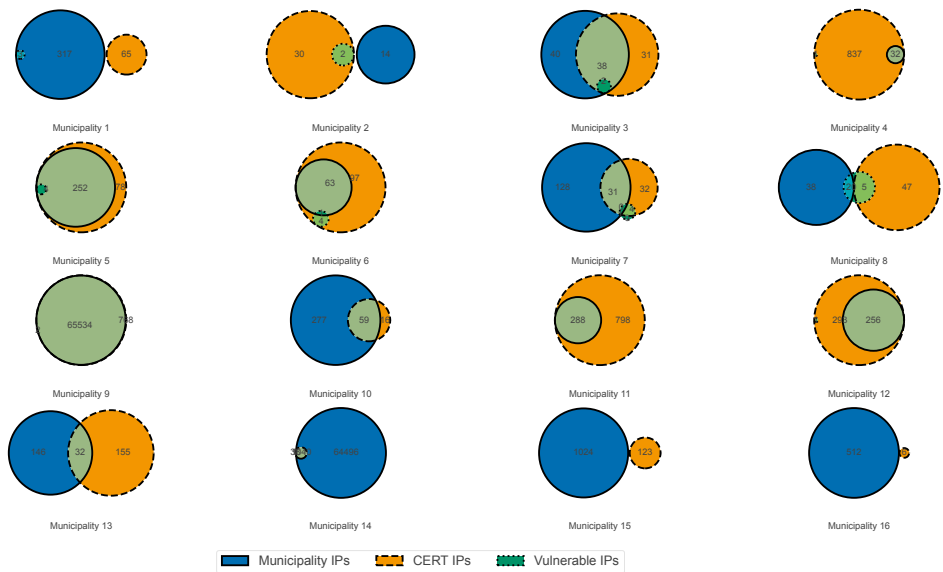


Figure A.1: IP sets registered with the CERT and the IPs used by a municipality.

# B

## APPENDIX FOR CHAPTER 3

### B.1. INTERVIEW PROTOCOL

The interview protocol followed a general structure across all interviews. Note that we asked participants to articulate a list of CSIRT services. The section “Experiences per Service X”, was repeated for each articulated service.

Minor variations in which questions were asked were introduced depending on the stakeholder group: governance, CSIRT staff, and constituents. It also presents the generic protocol and indicates which stakeholder groups received each question. A value of "Yes" means the question was asked of that group. "No" means the question was not asked. Some questions were slightly rephrased depending on the stakeholder group.

The interview protocol for other sector CSIRT practitioners, depicted in Table B.2, is almost identical, except for the organization name reference.

### CODES

**Table B.1:** Interview protocol with stakeholder-specific variations

<b>Section</b>	<b>Question</b>	<b>Const. CSIRT Gov.</b>		
<b>Introduction</b>	Can you tell me about yourself and your role?	✓	✓	✓
	What does a typical day look like?	✓	✓	✓
<b>Expectations of IBD</b>	In what capacity are you dealing with the IBD?	✓	✗	✓
	How would you describe that interaction?	✓	✗	✓
	What are your/the expectations of the services provided by the IBD?	✓	✓	✓
	What services are you using/providing, and has that changed over time?	✓	✓	✓
<b>Experiences per Service X</b>	For service X, how is it used?	✓	✓	✗
	For service X, what challenges are you facing?	✓	✓	✓
	For service X, has that changed over time?	✓	✓	✗
<b>Missing Services</b>	What services are you missing?	✓	✗	✗
<b>Outro</b>	What didn't we ask that we should have asked?	✓	✓	✓

"Gov." = governance stakeholders; "Const." = constituents; "CSIRT" = IBD staff.

**Table B.2:** Interview protocol for practitioners at other sector CSIRTs

<b>Section</b>	<b>Question</b>
<b>Introduction</b>	Can you tell me about yourself and your role?
	What does a typical day look like?
<b>Expectations of {ORG}</b>	What do you think are the expectations of the services provided by {ORG}?
	What services are you providing, and has that changed over time?
<b>Experiences per Service X</b>	For service X, how is it used?
	For service X, what challenges are you facing?
	For service X, has that changed over time?
<b>Outro</b>	What didn't we ask that we should have asked?

{ORG} denotes the specific CSIRT organization, depending on interviewee context.

**Table B.3:** Full Codebook (Services — Incident Response and Advisories)

Group	Theme	Subcode	Description
Services	Incident Response	Incidents	Participants explaining incidents they did or did not encounter and the value of the CSIRT during
		Boots-on-the-ground	Participants explaining their expectations of the CSIRT during an incident
		Responsibilities	Participants explaining who does what during an incident
		Communication	Participants explaining how they communicate about this service
		Technical Capabilities and Resources	Participants explaining the expected and actual available technical capabilities and resources to handle incidents
	Advisories	Incident Reporting and Regulation	Participants explaining if and how they report incidents to the CSIRT, and the regulatory context (NIS2 and BIO) for these processes
		Acting on Advisories	Participants describing the value of advisories and how they are acted upon
		Software Inventory and Reporting	Participants explaining how the inventories of their software are managed and reported to the CSIRT.
		Timeliness	Participants describing the timeliness of advisories
		Frequency	Participants describing how often they receive advisories
	False Positives	Participants describing the problem of false positives in whether or not they are running vulnerable software	
	Internal Decision-making	Participants describing the role of advisories in internal decision-making processes	

**Table B.4:** Full Codebook (Services — Expert Insights and Vulnerability Notifications)

Group	Theme	Subcode	Description
Services	Expert Insights	Operationalization Issues	Participants describing their difficulties in providing and maintaining templates of knowledge products
		Recipient Diversity	Participants describing the difficulties of tailoring the level of detail of products for mixed maturity levels of constituents
		Quality	Participants describing the value and quality of the service
		Legal Issues	Participants explaining legal issues in offering templates for certain processes
		Implementation Issues	Participants describing the difficulties in implementing the templates into actual processes within their organization
	Vulnerability Notifications	Asset Inventory	Participants describing the state of their current asset inventory
		Reporting Assets	Participants explaining the reporting process of their asset inventory to the CSIRT
		Frequency	Participants describing the value and frequency of vulnerability notifications
		Process	Participants describing their process to collect and update their asset inventory
		Responsible Disclosures	Participants describing the value, difficulties, and process of handling incoming responsible disclosures

**Table B.5:** Full Codebook (Services — Intelligence Sharing and Outreach)

Group	Theme	Subcode	Description
Services	Intelligence Sharing	Types	Participants describing the value of different kinds of intelligence
		Frequency	Participants describing that they often received shared intelligence
		Visibility	Participants describing that intelligence puts the CSIRT on their radar as a supportive organization
		Reputation	Participants describing the reputation of the CSIRT or fearing their own reputation in sharing back
	Outreach and Community	Knowing Constituents	Participants describing the value of personal relations and the importance of knowing the constituency and the CSIRT staff
		Building Trust	Participants describing the role of this service to increase trust with the CSIRT and among constituents
		Community Facilitation	Participants describing efforts and initiatives to bring constituents together

**Table B.6:** Full Codebook (Practitioner Organizational Challenges)

Group	Theme	Subcode	Description
Practitioner Organizational Challenges	Governance structure and stakeholders	Mandate and structure	Practitioners describing the CSIRT mandate and/or their organizational structure
		Organizational dependencies	Practitioners describing their interactions and/or dependencies on other organizations in providing their services
		Trust	Participants describing the value and necessity of trust to share information
		Constituent population	Participants describing their constituent population
	Infrastructure and capability management	Infrastructure and asset management	Participants describing their infrastructure and related (management) processes.
		Internal and external security services	Participants describing their security services and related processes

# C

## APPENDIX FOR CHAPTER 4

### C.1. CHARTS AND TABLES

**Table C.1:** Actor counts in public mappings

<b>Source</b>	<b># Actors</b>
MISP TAG	855
BreachHQ	760
Thai CERT	504
MITRE	171
Qianxin	66

**Table C.2:** Filetypes with only agreement

File type	Agree
ZIP	63
JavaScript	35
PDF	22
MS Word Document	14
Email	9
unknown	6
Rich Text Format	5
RAR	3
Mach-O	2
Text	2
DOS batch file	1
PGP Security Key	1
Shell script	1
Office Open XML Presentation	1
Windows Installer	1
INI	1
Compiled HTML Help	1

**Table C.3:** Filetypes with only disagreement

File type	Disagree
Office Open XML Spreadsheet	1

**Table C.4:** Ambiguous actor names in TAG ( $n = 23$ )

Name	Name	Name
circle typhoon	evasive panda	andariel
bronze highland	mint sandstorm	sapphire sleet
nobelium	pla navy	g0013
allanite	apt43	greenbug
stardust chollima	iridium	spandex tempest
golden chickens	imperial kitten	smoke sandstorm
cobalt gypsy	cadet blizzard	oilrig
grizzly steppe	cuboid sandstorm	




Figure C.1: Overlap of tracked actors between vendors.

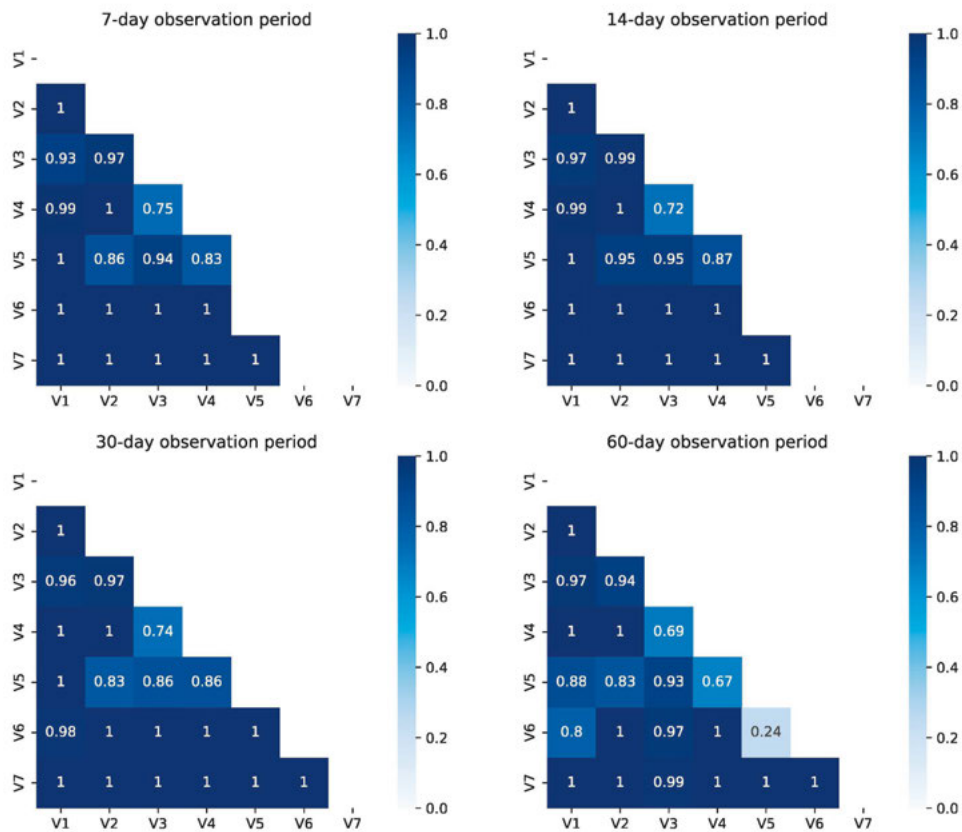


Figure C.2: Pairwise vendor agreement on country attribution for co-observed IOCs.

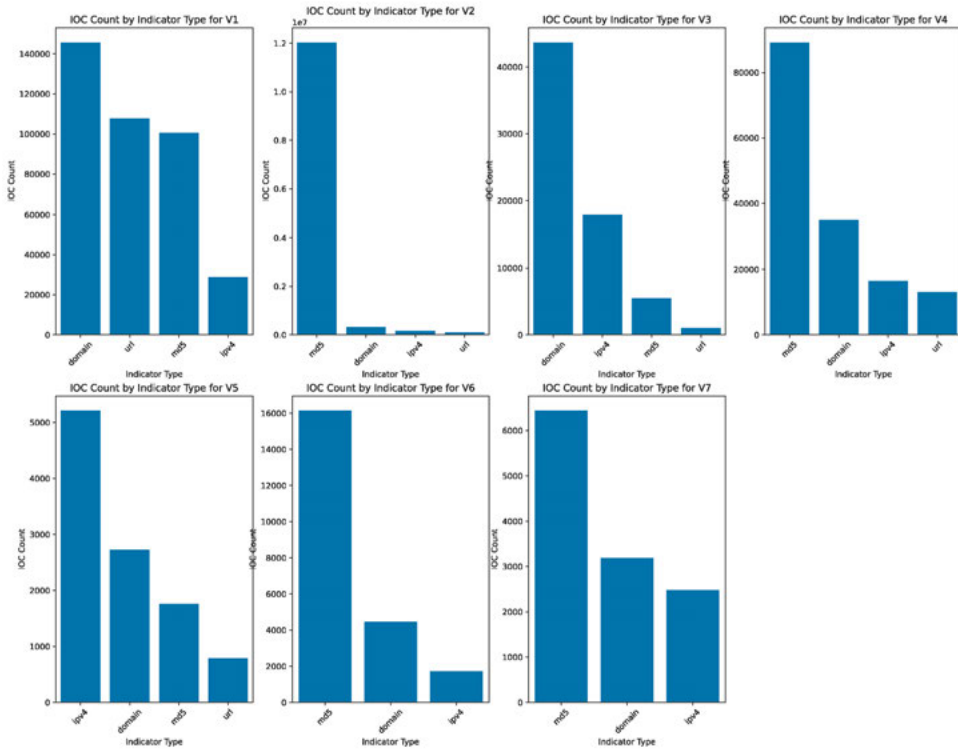


Figure C.3: IOC count per indicator type per vendor

**Table C.5:** CS/MS mapping comparison against TAG

<b>A. Actors not present in TAG</b>	
Amethyst Rain	VolcanicTimber
Highground	DRAGNET PANDA
Houndstooth Typhoon	CHATTY SPIDER
Storm-0252	MintedSoil (CyberRoot)
Wisteria Tsunami	RENEGADE JACKAL
Pinstripe Lightning	
<b>B. Names deconflicted in CS/MS but mapped to different actors in TAG</b>	
LABYRINTH CHOLLIMA	Jade Sleet
Midnight Blizzard	INDRIK SPIDER
Mustard Tempest	VELVET CHOLLIMA
Opal Sleet	LOTUS PANDA
Raspberry Typhoon	VELVET CHOLLIMA
Ruby Sleet	VICE SPIDER
Vanilla Tempest	ALPHA SPIDER
Velvet Tempest	Cozy Bear

**Table C.6:** Actor and country attribution agreement for 4 vendors. Values show Krippendorff's  $\alpha$  with 95% CI and sample size ( $n$ ).

Type	OP	4 vendors	
		Actor	Country
ALL	30	1.00 [1.00,1.00] ( $n=1$ )	1.00 [1.00,1.00] ( $n=1$ )
ALL	60	1.00 [1.00,1.00] ( $n=6$ )	1.00 [1.00,1.00] ( $n=6$ )
ALL (inc. WIP)	14	-0.17 [-0.17,0.00] ( $n=2$ )	-0.17 [-0.17,0.00] ( $n=2$ )
ALL (inc. WIP)	30	-0.06 [-0.24,0.07] ( $n=9$ )	-0.12 [-0.29,0.01] ( $n=9$ )
ALL (inc. WIP)	60	0.14 [-0.04,0.26] ( $n=20$ )	0.07 [-0.11,0.21] ( $n=20$ )
MD5	30	1.00 [1.00,1.00] ( $n=1$ )	1.00 [1.00,1.00] ( $n=1$ )
MD5	60	1.00 [1.00,1.00] ( $n=5$ )	1.00 [1.00,1.00] ( $n=5$ )
DOMAIN	60	1.00 [1.00,1.00] ( $n=1$ )	1.00 [1.00,1.00] ( $n=1$ )
URL	-		
IPV4	-		

**Table C.7:** Actor and country attribution agreement for ALL indicators (ex. WIP) for secondary indicator types. Values show Krippendorff's  $\alpha$  with 95% CI and sample size ( $n$ ).

Type	OP (Days)	2 vendors	
		Actor	Country
ALL	7	0.40 [0.33,0.46] ( $n=93$ )	1.00 [1.00,1.00] ( $n=93$ )
ALL	14	0.37 [0.32,0.42] ( $n=152$ )	1.00 [1.00,1.00] ( $n=152$ )
ALL	30	0.34 [0.30,0.38] ( $n=238$ )	1.00 [1.00,1.00] ( $n=236$ )
ALL	60	0.29 [0.26,0.32] ( $n=375$ )	1.00 [1.00,1.00] ( $n=372$ )

# ACKNOWLEDGEMENTS

Before embarking on my PhD, I imagined it would be a solitary and lonely journey. That turned out to be quite wrong. Instead, I discovered it was really a journey traveled with the help of many others. Rather than going at it alone, it was very much a team effort. This section is a tribute to those people who were part of my PhD journey.

Before diving into the gratitudes, I want to highlight one comment that my daily supervisor, Rolf, once told me that has stayed with me ever since. He said that “every PhD candidate has their own unique skills *and* luggage”. Having accumulated a fair amount of professional experience before starting my PhD, I often felt considerable pressure to perform well. That remark brought me comfort and gradually helped me understand that a PhD is, for everyone, a path of growth, and therefore also one that comes with at least some degree of suffering, frustration, and despair, regardless of your background or experience. Everyone has their own growing pains I guess. In the end, it’s simply a journey in which you need to persist, and one learns that the sweet is never as sweet without the sour. I am glad I persisted, but that would not have been possible without the help of many others who pulled me through. As William Shakespeare once wrote, *“I can no other answer make but thanks, and thanks, and ever thanks.”*

To my promotor Michel, first of all, thank you for taking me on board as a bit of an odd PhD candidate. I wasn’t quite sure that a PhD was for me, but our earliest mail exchanges instantly gave me a feeling of comfort, and a sense that this route was the way forward. You also set a valuable example of how to provide a safe space for the exchange of ideas and delivering feedback. People and ideas were always taken seriously while there was always a place for fun and laughter. This doesn’t mean that debates could not be heated or ideas downright killed, but always with respect and (analytical) integrity. I think many (academic) environments could learn from how you run the show.

I also deeply respect your analytical prowess. It’s amazing to see how quickly you can grasp vague ideas, mainly by myself, and articulate them so clearly. Or sometimes take one of these vague ideas and build on them so quickly, freely, and creatively. I hope one day you get discovered by the mainstream media and start providing comments in late night TV shows on whatever topic happens to be relevant that day. It’s time van Rossem makes room! Oh yes, by the way, given your speed of thought, your patience with PhD’s is also highly admirable.

To Rolf, thank you for many things, including listening. I really appreciate our coffee sessions where we could talk and you would always listen to how things were going. I also find your enthusiastic “can-do” and hustle attitude an inspiration. To bend the rules, one needs to know the rules precisely, and this is something that you taught me. Additionally, I have fond memories of hanging around in your office: the hipster social gathering place of TBM. Your coffee machine was the place to be to learn about rumors, share a joke, or just engage in random banter. And most importantly, you taught me the basics of coffee snobbery at that machine *and* during our trip to New York. I will make sure to hone those skills. Speaking of that office, I also want to give a big shoutout to Joyce, I really enjoyed hanging out with you and talking (nonsense), especially as a fellow parent. As an all-round fun and nice person, people are blessed with your presence.

To Yury, thank you for being such a kind, supportive, and patient person. Your door has always felt open and you have always been very accessible to help. Your kindness and support are an inspiration. I vividly remember sharing my first colloquium presentation, which was filled with

images, misaligned text, quotes and whatnot, and you said something like: "this is great, this is *you*. You just need to articulate the research question in a slide also". Moreover, your technical skill is unparalleled and it was a blessing to have you available for technical assistance. You could spot and fix bugs or optimizations in an instant. For example, I remember an IP matching script that took hours and, after one glance, you gave a tip to optimize it to minutes. Thanks for all the support.

To Remko, Bas, and the others from the Informatiebeveiligingsdienst (IBD), thank you for listening to my research proposal early on, even when the research may have been a bit vague. I also appreciate the amount of trust and level of support and access that was provided by the IBD. The amount of reflection and vulnerability that your team and your organization have shown is an example for many other organizations in the Dutch cyber ecosystem. I wish you all the best of luck in your new endeavors.

To Ingmar, and the Ministry of Interior, thank you for working with this irregular PhD candidate and guiding me around the world of civil servants. I appreciate our coffees and chats together and your help in connecting me with the right people. It is an honor to have you on my committee.

To the other external collaborators, in particular the people at the NCSC, but also Frank, Erik, Jennie, Stef, Bram, Niels, and many more, you know who you are! Thank you for collaborating with me and supporting me. I have always felt accepted, seen, and heard by all of you. You helped me succeed in finishing the PhD. Thank you for your support.

To all the participants, thank you for taking my research seriously and participating, sharing on-record your candid and honest views on your experiences in the Dutch cybersecurity world. Your honesty and vulnerability are admirable and will make the cybersecurity world better.

To Natalia, thank you for taking care of me when I first arrived in Delft. You quickly became a colleague and a friend and I value both very much. There are so many fond memories already: from drinking games to conducting interviews to karaoke and cosplay events. Your social efforts have been instrumental in forming the kind, tight-knit group at TBM that we are today. Even if it means that one sometimes has to suffer awkward karaoke or drink unnamed (strong) drinks. I hope life treats you and your family well.

To Xander, thank you for showing me around PhD life, for many fun chats, and for introducing me to so many interesting people. You are a kind and smart person that wishes others the best, a very admirable quality. I'm pretty sure we will run into each other in the future. I'm convinced you have a bright future ahead. And good luck with those twins!

To Yana, thank you for teaching me things on so many levels. First, as a master student you taught me how to supervise. It really helped to work with a formidable student like you. Then, as a PhD colleague, you showed me how to work hard, dedicated, and with determination. I am sure you will produce great research and make many friends along the way.

To my office mates, Lorenz, Ronak, and Evi. Lorenz, we initially connected on Star Wars and gaming, but soon we started hanging around often at the office and chatting about whatever topic. At the start of our PhD we had dinner in Amsterdam, and hey, look at us now, we both made it! You became a wonderful father during your PhD, managing your time effectively, and still producing top-tier research, amazing. Along the way you even became a bicycle snob! I wish you all the best buddy. Ronak, I will fondly remember our office chats and coffees. You are an inspiration in looking at life positively and with kindness. You have grown so much during your PhD, it was an honor to get to know you and to watch your growth so closely. Evi, we chatted so many times in the office, and I loved it every time. I vividly remember when one of us would come back after a supervision meeting and vent on how nervous we were or how we could not articulate our ideas clearly. Additionally, we must have discussed the entire set of Nintendo Switch games, with a particular focus on cosy gaming. One day you will have a great digital farm.

Simon, thank you for helping me with the CHI paper. Your insights and feedback were in-

valuable in making it a success. On a personal level, I love how we could chat about music and instruments. Maybe one day I will see you slapping that bass. You, too, are filled with kindness and intelligence. Our department is lucky to have you and I hope you climb the ranks accordingly, you deserve it.

Savvas, man, you have often made me laugh. I love how your research interest, memes, has somehow also transferred to the physical world. Your real-world interactions in the group, in some way, always felt "meme-y". A remark like "gosh" is forever audibly ingrained in my brain, as is "bro" and "why so toxic". You are incredibly gifted as a researcher and I hope your career will reflect it. By the way, I was waiting at the entrance on Friday at seven but you were not there bro, what gives dude?

Carlos, fellow dad, our interactions professionally were limited but we had many casual chats in your office. I was always impressed with how friendly and funny you are in one moment, and the next moment you can switch to some difficult scientific method. It was also very nice to share early experiences with a fellow dad in the trenches. Remarkably, thanks to you, I also improved my knowledge and opinions on toilets along the way, how odd.

To my fellow PhD candidates! Elsa, thank you for taking such good care of me and the other fellow PhD's. You always made sure nobody was left behind. Szu, thank you for being so kind to my kids and myself. Ezra is still amazed by the kaleidoscope you gave him. Veerle, thanks for the many parenting tips and (vegan) recipes. I'm confident I'm a healthier person thanks to you. Mathew, thanks for helping me with getting data, fine-tuning research ideas, and being an all-round nice guy. Similar to Elsa, you always look after others, bless you. Hugo, thanks for listening to my endless questions on the administrative side of finishing a PhD. You are a fun and all-round cool dude who, surprisingly, *can* compare apples with oranges! I hope we stay in touch. Swaathi, thanks for showing so much dedication to research. I hope your academic journey continues to soar. Pepijn, our interactions were limited, but you too are a friendly and smart guy. I hope you thrive in your PhD! Max, my fellow watch- and bike enthusiast, I hope PhD life continues to treat you well. Maybe one day we will grill a steak. Radu, thanks for helping me set up a nifty home network. I run a neat home lab thanks to you. Kelvin, thanks for being around, always kind, helping and listening. You are super smart and I'm convinced that, one way or the other, you will be critical in some big master plan to save the world. Cecile, thank you for always checking in with me and the others, and helping me make sure to take myself and my feelings seriously. Fieke, thanks for showing me how to combine a busy professional life with an academic journey. You work incredibly hard and it's impressive. Arwa, thanks for the endless chats on dogs. I hope one day you will be flooded by a sea of puppies. Annel, thanks for introducing me to *real* hardcore music. I had to get used to it a bit, but you opened my mind to new music. You also had many "good vibes". Sandra, no matter how depressed one is, your smile and laughter provide instant relief and joy. I deeply respect how you have grown as a researcher and now combine both parenting life and academia. I hope you are happy and continue to brighten many people's lives. Gebrand, fellow old guy of the group, you have always approached everything with an open mind and with kindness, I hope life treats you well. To Maaïke, thanks for showing me the importance of physical fitness. I'm impressed with how you combine a physically disciplined lifestyle with so much hard academic and professional work. I also love the fact that we connected well on so much nerdy stuff. To Marie, thanks for being an inspiration in undertaking so many different activities. Another thing I will remember, is that whenever I was in your office, somehow, there was also always some candy! To Abraham, fellow coder, I love how you made a case for pumpkin spice coffee. It's a bold and brave move, similar to having pineapple on pizza. I hope you sail smoothly throughout the rest of your PhD. To Vahid and Nicolo, we didn't interact all that much. However, I got a taste of your potential during a karaoke session at Michel's. Given that performance, I foresee a bright future for the both of you. Take care of those x-wing's and good luck with the PhD!

To the other supportive people at Delft, Joyce, Joy, Nicolas, and the kind lady of the lunch shop, thank you for being part of my PhD life and helping me along the way. We had many kind chats and interactions and each of you have supported me in one way or another. I wish you all the best. To anyone else at Delft who I forgot, thank you for your support!

Papa, mama and Alena, thank you for always supporting me and believing in me. I cannot even begin to express my gratitude for a lifetime of support and love in a few sentences, but you have raised me into the person I am today and I'm very proud of that. Thank you and I love you.

Michelle, my love, thank you for being there for me the entire time. You stuck by me from the moment I started right to the finish. Somebody once told us that a PhD is like having a 'third person' on the couch. And boy, that was very accurate. Soon that third person will be gone and we will have plenty of space for both of us and our lovely two sons! I could not have done it without you and every day I am blessed to have you as my girlfriend, wife, and mother of our children. I love you incredibly much and I strongly believe that you should have some kind of formal academic title for sticking with a PhD'er too.

Ezra and Elias, my lovely sons, I am so incredibly proud of you two! You were both born during my PhD and you are the best things in my life. I hope each of you finds inspiration in the kindness, support, honesty, vulnerability, and patience that the people around me during my PhD have shown me. I love you both incredibly much.

*Aksel Ethembabaoglu  
Rijswijk, February 2026*

# AUTHORSHIP CONTRIBUTIONS

This dissertation is founded on three peer-reviewed papers deriving from studies conducted with several co-authors. Each study, and this dissertation, truly is the result of teamwork. While I am first author on each study, I benefited greatly of the insights and feedback of my collaborators. In the next paragraphs I will outline the specific contributions of each co-author for every study.

For the first study (Chapter 2), Yury Zhauniarovich, Rolf van Wegberg and Michel van Eeten helped scope the study, structure the research, writing the draft, proofreading and polishing of the text. I conducted the collection and analysis of data and did most of the writing.

For the second study (Chapter 3), Natalia Kadenko helped arrange interviews, and conduct and analyze the interviews. Yana Angelova helped building the data pipeline for the analysis and collection of vulnerability notifications. Simon Parkin helped in framing the study from a human-factor security perspective, he also suggested additional literature and reviewed the earlier research structure and writing. Similar to the previous study, Yury Zhauniarovich, Rolf van Wegberg and Michel van Eeten helped scope the study, structure the research, writing the draft, proofreading and polishing of the text. I conducted the interviews and did the analysis of data. I also did most of the writing.

For the third study (Chapter 4), Yury Zhauniarovich, Rolf van Wegberg and Michel van Eeten helped scope the study, structure the research, writing the draft, proofreading and polishing of the text. Yury, in particular, made a strong pass at the Introduction. I conducted the data collection and analysis of the study. I also did most of the writing.

Finally, I want to highlight the great efforts of Michel van Eeten in supporting me to write high-quality papers.



# LIST OF PUBLICATIONS

5. **Ethembaoglu, A.M.**, van Wegberg, R.S. & Zhauniarovich, Y. & van Eeten, M.J.G. (2024). “The Unpatchables: Why Municipalities Persist in Running Vulnerable Hosts”. In *Proceedings of the 33rd USENIX Security Symposium (USENIX Security '24)*.
4. **Ethembaoglu, A.M.**, Kadenko, N.I., & Angelova, Y., & Zhauniarovich, Y. & Parkin, S. & van Eeten, M.J.G. (2026). ““Tell Them They Are a Responsible Entity, Not a Customer”: Understanding Practitioner Challenges in Sector CSIRTs”. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*
3. **Ethembaoglu, A.M.**, van Wegberg, R.S. & Zhauniarovich, Y. & van Eeten, M.J.G. (2026). “APT to Disagree: A Comparative Analysis of Attribution in Commercial TI”. In *Proceedings of the IEEE Symposium on Security and Privacy (S&P '26)*.
2. Bouwman, X., **Ethembaoglu, A.M.**, Hermans, B. & Gañán, C. & van Eeten, M.J.G. (2025). “Can IOCs Impose Cost? The Effects of Publishing Threat Intelligence”. In *Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security (CCS '25)*.
1. Angelova, Y., **Ethembaoglu, A.M.**, van Wegberg, R.S. & Gañán, C. & van Eeten, M.J.G. (2026). “I’m the mess that you wanted: Evaluating the accuracy of WHOIS asset discovery against self-reported data”. In *Proceedings of the 25th Workshop on the Economics of Information Security (WEIS '26)*.



# ABOUT THE AUTHOR



Aksel Ethembabaoglu (1982) was born in The Hague, the Netherlands. In 2021, he joined Delft University of Technology as a PhD candidate. Previously, he received his BSc degree in Artificial Intelligence at the University of Amsterdam and his MA degree in International Relations at King's College London. Aksel has worked at the intersection of technology, geopolitics, and security in various capacities in government and industry.

During his PhD, he focused on the human and organizational aspects of cybersecurity in organizations, using socio-technical measurements to gain relevant new insights that improve governance. He was also involved in assisting in cybersecurity courses and supervising master students.

His research interests in cybersecurity include: threat intelligence, human factors, vulnerability management, CSIRTs, and attackers (APTs). He is a firm believer in interdisciplinary research.