THE IMPACT OF GOOD PRACTICES ON THE SECURITY OF DIGITAL PRODUCTS AND THE LIMITATION OF OFFENSIVE CYBER OPERATIONS

By Vladimir Radunovic

WORD COUNT 44344 TIME SUBMITTED 06-DEC-2025 03:57PM
PAPER ID 119334112

SINGIDUNUM UNIVERSITY
BELGRADE

DEPARTMENT FOR POSTGRADUATE STUDIES

DOCTORAL THESIS

THE IMPACT OF GOOD PRACTICES ON THE SECURITY OF DIGITAL PRODUCTS AND THE LIMITATION OF OFFENSIVE CYBER OPERATIONS

Mentor
Prof. Aleksandar Jevremović

Student
Vladimir Radunović
Student ID: 460062/2012

BELGRADE, 2025

Mentor:

Prof. Aleksandar Jevremović, Singidunum University, Belgrade, Serbia

Committee:

1. Prof. Aleksandar Jevremović, Singidunum University, Belgrade, Serbia
2. Prof. Mladen Veinović, Singidunum University, Belgrade, Serbia
3. Prof. Ivan Čorbev, Ss. Cyril and Methodius University, Skopje, North Macedonia

Date:

Acknowledgments

To my unbelievably patient family.

To Jovan Kurbalija and colleagues at Diplo, who led me into this journey.

To professors Mladen Veinović and Aleksandar Jevremović, and the University staff, who never gave up on me.

“Everything will be okay in the end. If it's not okay, it's not the end.”
John Lennon

Abstract

Title: The Impact of Good Practices on the Security of Digital Products and the Limitation of Offensive Cyber Operations

Abstract:

Vulnerabilities in digital products are a critical enabler of offensive cyber operations. States increasingly procure and exploit vulnerabilities, contributing to cyber armament and systemic insecurity in cyberspace. Alongside diplomatic, normative, and regulatory initiatives, an important pathway to constraining this trend is reducing vulnerabilities through security-by-design and vulnerability management practices. While leading technology suppliers have developed a range of such practices – including Product Security Incident Response Teams (PSIRT) as a unit responsible for managing vulnerabilities – there remains a lack of empirical methodologies for assessing whether these practices measurably reduce vulnerability exploitation in real-world scenarios.

This thesis investigates the effectiveness of industry security-by-design practices – with a focus on PSIRT – in reducing the likelihood of active vulnerability exploitation and explores the implications for offensive cyber operations. Drawing on a structured review of academic literature, standards, political and regulatory approaches, and industry practices, it discusses the reality of cyber armament and the related role of vulnerability exploitation, analyses industry practices for vulnerability management, and identifies gaps in existing assessment methods for their real-world impact. In response, it proposes a novel and reusable methodology for evaluating the effectiveness of security-by-design practices, based on a randomised matched case-control design.

Applying this methodology to publicly available data demonstrates that the presence of a PSIRT within a supplier is associated with a statistically significant absolute risk reduction of 17% in exploitation likelihood. Through a broader interpretation of results, the study concludes that PSIRT constitutes a meaningful security-by-design practice that measurably reduces the likelihood of active vulnerability exploitation and, when embedded within a broader ecosystem of secure development, regulatory and political measures, increases the cost, uncertainty, and operational risk of exploit-based attacks. It thereby supports – but does not alone determine – the potential of industry practices to constrain sophisticated offensive cyber operations. The proposed methodology enables future empirical research on security practices and evidence-based evaluation of their security and geopolitical impact.

Keywords: vulnerabilities, security-by-design, security of digital products, secure development lifecycle, cyberattacks, offensive cyber capabilities, offensive cyber operations

Sažetak

Naziv: Uticaj dobrih praksi na bezbednost digitalnih proizvoda i ograničavanje ofanzivnih sajber-operacija

Sažetak:

Ranjivosti u digitalnim proizvodima predstavljaju ključni faktor za izvođenje ofanzivnih sajber-operacija. Države sve češće pribavljaju i zloupotrebljavaju ranjivosti, doprinoseći sajber-naoružavanju i sistemskoj nesigurnosti u sajber-prostoru. Uz diplomatske, normativne i regulatorne inicijative, jedan od važnih načina za ograničavanje ovog trenda jeste umanjenje ranjivosti kroz pristupe bezbednosti kroz dizajn (*security-by-design*) i prakse upravljanja ranjivostima. Iako su vodeći tehnološki proizvođači razvili niz takvih praksi – uključujući uspostavljanje timova za odgovor na bezbednosne incidente proizvoda (*Product Security Incident Response Team* – PSIRT) kao specijalizovanih jedinica za upravljanje ranjivostima – i dalje postoji nedostatak empirijskih metodologija za procenu da li ove prakse merljivo smanjuju zloupotrebu ranjivosti u realnim uslovima.

Ova disertacija ispituje efikasnost industrijskih praksi bezbednosti kroz dizajn – sa posebnim fokusom na PSIRT – u smanjenju verovatnoće aktivne eksploatacije ranjivosti i razmatra njihove implikacije na ofanzivne sajber-operacije. Polazeći od strukturiranog pregleda akademske literature, standarda, političkih i regulatornih pristupa i industrijskih praksi, rad razmatra realnost sajber-naoružavanja i ulogu eksploatacije ranjivosti u tome, analizira industrijske prakse za upravljanje ranjivostima i identifikuje nedostatke postojećih metoda za procenu njihovog uticaja u praksi. Kao odgovor na te nedostatke, predlaže se nova i ponovljivo primenljiva metodologija za evaluaciju efikasnosti praksi bezbednosti kroz dizajn, zasnovana na randomizovanoj studiji istraživanja parova (*randomised matched case-control design*).

Primena ove metodologije na javno dostupnim podacima pokazuje da je prisustvo PSIRT-a kod proizvođača u korelaciji sa statistički značajnim apsolutnim smanjenjem rizika eksploatacije od 17%. Kroz šire tumačenje rezultata, rad zaključuje da PSIRT predstavlja značajnu praksu bezbednosti kroz dizajn koja merljivo smanjuje verovatnoću aktivne eksploatacije ranjivosti i, kada je uklopljena u širi ekosistem bezbednog razvoja, regulatornih i političkih mera, povećava troškove, neizvesnost i operativni rizik eksploatacionih napada. Time se potvrđuje – ali se ne određuje isključivo – potencijal praksi kompanija da doprinesu ograničavanju sofisticiranih ofanzivnih sajber-operacija. Predložena metodologija omogućava buduća empirijska istraživanja bezbednosnih praksi i procenu njihovog bezbednosnog i geopolitičkog uticaja zasnovanu na činjenicama.

Ključne reči: ranjivosti, bezbednost kroz dizajn, bezbednost digitalnih proizoda, bezbedni ciklus razvoja, sajber napadi, ofanzivni sajber-kapaciteti, ofanzivne sajber-operacije

Table of Contents

Acknowledgments 1
Abstract 2
Sažetak 3
Table of Contents 4
Abbreviations 6
Figures and Tables 7
Figures 7
Tables 7
1 Introduction 8
1.1 Background 8
1.2 Research Subject 10
1.3 Objectives 10
1.4 Hypothesis 11
1.5 Research Questions 12
1.6 Scientific Contribution 12
1.7 Structure of the Thesis 13
2 Review of the Research Field 15
2.1 Growing Cybersecurity Risks 16
2.2 Increasing the Costs of Cyber Armament and Vulnerability Exploitation 28
3 Industry Models and Practices for Reducing the Vulnerabilities in Digital Products 39
3.1 Models and Methods from the Academic Sector 39
3.2 International and National Standards and Recommendations 41
3.3 Secure Development Lifecycle Practices of Leading Suppliers 43
3.4 Vulnerability Management and Coordinated Vulnerability Disclosure 45
4 Assessing the Efficiency of Good Practices 59
4.1 Industry Methodologies 59
4.2 Academic Methodologies 61
4.3 Key Elements for Assessing PSIRT Efficiency 63
5 Proposed Methodology for Assessing the Efficiency of PSIRT 67
5.1 Matched Case-control Design 67
5.2 Variables 69
5.3 Datasets 73
5.4 Data Pre-processing 76
6 Model Deployment and Calculation 78
6.1 Sampling, Matching and Bootstrapping 78

6.2 Risk Reduction 80
6.3 Testing Statistical Significance 82
7 Results 84
7.1 Risk Reduction 84
7.2 Statistical Significance 86
8 Discussion 88
8.1 A Mature but Uneven Landscape of Industry Practices 88
8.2 Association Between PSIRT Adoption and Lower Exploitation Risk 90
8.2 Impact of PSIRT on Offensive Cyber Operations 92
8.3 Impact of Industry Security Practices on Offensive Cyber Operations 94
9 Conclusion 97
Literature 100
Dataset References 118

Abbreviations

AI – Artificial Intelligence
APT – Advanced Persistent Threat
ASEAN – Association of Southeast Asian Nations
CERT – Computer Emergency Response Teams
CI/CD – Continuous Integration and Continuous Delivery/Deployment
CI – Critical (societal) infrastructure
CPE – Common Platform Enumeration
CVD – Coordinated Vulnerability Disclosure
CVE – Common Vulnerabilities and Exposures
CVSS – Common Vulnerability Scoring System
DAST – Dynamic Application Security Testing
DevSecOps – Development, Security, and Operations
EDB – Exploit-db
EoL – End-of-life (of the product)
FIRST – Forum of Incident Response and Security Teams
IoT – Internet of Things
IoMT – Internet of Medical Things
ICS – Industrial Control Systems
KEV – Known Exploited Vulnerability (database maintained by US CISA)
NVD – National Vulnerability Database (database maintained by US NIST)
OAS – Organization of American States (OAS)
OCC – Offensive cyber capabilities
OCO – Offensive cyber operations
OSCE – Organization for Security and Co-operation in Europe OSS – Open-Source Software
OT – Operational Technologies
PLC – Programmable Logic Controllers
POC – Proof-of-concept
PSIRT – Product Security Incident Response Team
SAST – Static Application Security Testing
SBOM – Software Bill of Materials
SCA – Software Composition Analysis
SCADA – Supervisory Control and Data Acquisition systems
SDL – Secure Development Lifecycle
SOC – Security Operation Centre
SSE – Secure Software Engineering
TPC – Third-Party Components
TTP – Tactics, techniques, and procedures
UN GGE – United Nations Group of Governmental Experts (on ICT Developments in the Context of International Security)
UN OEWG – United Nations Open-Ended Working Group (on Security of and in the Use of ICTs)

Figures and Tables

Figures

Figure 1: Map of states with offensive cyber capabilities 20
Figure 2: The Kaplan-Meier curve of the time-to-patch of vulnerabilities 27
Figure 3. Eleven UN norms for responsible state behaviour in cyberspace. 30
Figure 4: SDL mapped onto the standard digital product development process. 45
Figure 5. PSIRT role in a CVD process. 54
Figure 6: Geographical distribution of identified PSIRTs. 56
Figure 7: Distribution of PSIRTs by company size. 57
Figure 8. Distribution of PSIRTs by industry sector 58
Figure 9: PSIRT efficiency assessment based on matched case-control methodology 79
Figure 10: Median values of risk reduction parameters re. on bootstrap iterations 84
Figure 11. Median results related to risk reduction re. confounding variables 86

Tables

Table 1. Key International Standards and Frameworks for Digital Product Security 43
Table 2: Synthesis of the principal assessment elements identified across the literature 63
Table 3: Labelling of the observational units according to exposure and outcome status 68
Table 4: Labelling of the observational units according to PSIRT and the exploitation 70
Table 5: The overview of variables, types, and values they may take 72
Table 6: The number of vulnerabilities and suppliers in each pre-processed database 77
Table 7: Median results related to risk reduction with 200 iterations 84
Table 8: Median results related to statistical significance with 500 iterations 86

1 Introduction

Digital products currently underpin critical infrastructure, government functions, and global economic activity, while simultaneously embedding vulnerabilities that create systemic security risks. These weaknesses are no longer exploited only by cybercriminals but are increasingly harnessed by states as tools of power, placing product security at the intersection of technology, geopolitics, and international security.

This introduction chapter places the research within this interdisciplinary context, outlining the background and motivation for the study, defining its subject and objectives, and formulating the central hypotheses and research questions. It further sets out the scientific contribution of the work and provides an overview of the dissertation's structure.

1.1 Background

Digital products lie at the heart of contemporary digital transformation. Broadly defined, they include software, hardware, and hybrid systems capable of executing code, processing data, or communicating with other digital components (Kazakova, Radunović, and Droz 2023)[1]. In practice, digital products are embedded across an extensive range of applications, from standalone software and consumer devices to industrial control systems and cloud-based infrastructures. Their deployment underpins almost all critical societal and economic functions, supporting government operations, commercial services, healthcare, education, and essential infrastructure.

As reliance on digital products deepens, however, their security has become a growing source of systemic risk. Vulnerabilities in software and hardware introduce exploitable weaknesses that may be abused not only by opportunistic cybercriminals but also by highly resourced and strategically motivated adversaries. Such exploitation can compromise individual privacy, organisational integrity, and national security. The Organisation for Economic Co-operation and Development (OECD, 2022) warns that insecure digital products can produce cascading effects across interconnected systems, threatening the resilience of critical societal services.

The exploitation of vulnerabilities plays a central role in advanced persistent threats (APTs). These attacks – often attributed to state-sponsored or state-aligned actors – target high-value governmental and industrial assets, aiming at long-term access, intelligence collection, disruption, or coercion (Chen et al, 2014). Prominent analytical frameworks explicitly identify exploitation as a pivotal phase within the attack lifecycle. In the Lockheed Martin Cyber Kill Chain™, exploitation enables the installation of malicious payloads and subsequent command-

[1] Our research draws on insights and accumulated knowledge from the author's involvement in the Geneva Dialogue on Responsible Behaviour in Cyberspace, an international process established in 2018 to map the roles and responsibilities of key stakeholders – including states, the private sector, civil society, academia, and the technical community – in promoting security and stability in cyberspace. The Geneva Dialogue is led by the Swiss Federal Department of Foreign Affairs (FDFA) and implemented by DiploFoundation. See: https://genevadialogue.ch/

and-control (Hutchins et al., 2011). Likewise, the MITRE ATT&CK framework situates exploitation within the access phase common to most APT campaigns (MITRE n.d.). Similarly, all major APT models incorporate exploitation under the 'access' phase, affirming its strategic centrality (Lehto, 2022). Together, these models underline exploitation as a strategic enabler of sophisticated cyber operations.

At the state level, vulnerability exploitation has become a cornerstone of offensive cyber capabilities (OCC). In an environment of escalating geopolitical competition, states increasingly invest in tools, expertise, and markets that enable the discovery, acquisition, and retention of exploitable vulnerabilities – particularly so-called zero-day vulnerabilities, unknown to vendors and therefore lacking available patches. These capabilities are integrated into broader military, intelligence, and hybrid strategies. Their deployment through the offensive cyber operations (OCO), typically against foreign governmental or societal systems, risks undermining critical infrastructure, eroding trust in digital ecosystems, and contributing to the destabilisation of international security.

These developments highlight the urgent need to reduce both the availability and exploitability of vulnerabilities. One of the most direct and impactful ways is the improvement of how digital products are designed, developed, maintained, and supported throughout their lifecycle. This raises two interrelated questions: what responsibilities do suppliers[2] of digital products have in mitigating vulnerability-related risks, and how effective are their current practices in achieving this goal?

In response, industry and academia have developed a range of approaches aimed at strengthening secure development lifecycles (SDL)[3]. These include methodologies for secure software and hardware engineering, threat modelling, vulnerability management, coordinated disclosure, and patch deployment, as well as organisational mechanisms such as Product Security Incident Response Teams (PSIRTs). Collectively, these efforts form part of the broader security-by-design paradigm, which seeks to embed security considerations from the earliest stages of product conception rather than treating them as a post-deployment concern.

Although practices such as PSIRTs are widely recognised as good practice in managing product security, their real-world effectiveness remains insufficiently examined. Case studies and practitioner experience suggest that they improve coordination, accelerate remediation, and enhance communication with external stakeholders. However, there is limited empirical

[2] Academic and professional literature employs a range of terms – including *supplier*, *vendor*, *manufacturer*, and *(ICT) industry* – often with subtle differences in meaning. While distinctions between these terms may be relevant in specific regulatory or economic contexts, this work uses such terms interchangeably for clarity and consistency. This usage follows the OECD definition of *supply-side actors* as those "involved in the product or service's supply chain and having a role in developing, providing access to, maintaining, or selling the product or service, such as manufacturers, service providers, component providers, integrators, intermediaries and vendors, as well as third-party code owners". (OECD, 2022)

[3] Both secure development lifecycle (SDL) and secure software development lifecycle (SSDL) are the terms used in practice. Though SDL is also commonly used to mean the regular software development lifecycle, while SSDL is a more precise and common acronym for the secure version, the latter narrows the scope to software. Since our research addresses security of digital products, we adopt the SDL acronym throughout the work.

evidence demonstrating whether such practices systematically reduce the likelihood of vulnerabilities being exploited in the wild – the active exploitation in real-world scenarios, outside of a controlled or laboratory environment – or whether they constrain OCO.

Moreover, while maturity models and compliance frameworks help assess the procedural adoption of secure development practices, relatively few studies investigate the causal relationship between these practices and observed exploitation outcomes. This points to the need for analytical methods capable of linking organisational practices to measurable security effects. In particular, there is a need for methodologies that can assess whether security-by-design practices increase the complexity, cost, or uncertainty faced by adversaries seeking to exploit residual vulnerabilities.

This research addresses these gaps by proposing and implementing a novel empirical approach to evaluating the effectiveness of security-enhancing practices. It develops a methodological framework and applies it to one recommended practice: the establishment of a PSIRT. By examining how the presence of PSIRTs relates to patterns of vulnerability discovery, disclosure, patching, and real-world exploitation, the study assesses their impact on exploitation risk and, by extension, on the feasibility and effectiveness of OCO. In doing so, it aims to contribute to scholarly debates on cybersecurity risk reduction while also informing industry practice and policy discussions on countering the proliferation of OCC.

1.2 Research Subject

The subject of this research is the identification, classification, and analysis of existing and emerging best practices for the secure development of digital products, as well as the assessment of their effectiveness in managing vulnerabilities. The dissertation focuses in particular on practices aimed at managing vulnerabilities and reducing the risks associated with their exploitation throughout the product lifecycle.

Central to the analysis is the Product Security Incident Response Team (PSIRT) as an organisational practice for coordinating vulnerability discovery, assessment, disclosure, and remediation. The research proposes, and applies, a methodological model to verify the effectiveness of PSIRT in reducing the likelihood of real-world vulnerability exploitation.

Finally, the dissertation examines the broader implications of such security-by-design and vulnerability management practices for the feasibility of offensive cyber operations, the development and use of cyber capabilities, and the resulting effects on international security and geopolitical stability.

1.3 Objectives

The overall objective of this research is to assess whether and to what extent security-by-design practices adopted by suppliers of digital products – particularly the establishment of PSIRTs –

effectively reduce the risk of vulnerability exploitation in the wild and thereby increase the cost and complexity of conducting offensive cyber operations. This objective is pursued through the systematic identification and classification of leading industry practices, and through the development and empirical testing of a methodological model for evaluating their real-world effectiveness.

Theoretical objective:
To contribute to the interdisciplinary understanding of cyber armament by analysing how vulnerabilities in digital products enable offensive cyber operations and how security-by-design practices can alter this strategic risk environment.

Practical and technical objective:
To develop and apply an empirical methodology for assessing the effectiveness of PSIRTs and related vulnerability management practices in reducing the likelihood of active vulnerability exploitation in real-world conditions.

Societal objective:
To inform policymakers, regulators, and industry stakeholders on how corporate security practices – when systematically adopted – can complement diplomatic, regulatory, and normative efforts aimed at limiting cyber armament and enhancing overall international peace and stability.

1.4 Hypothesis

General hypothesis:
The geopolitical trend of cyber armament and the execution of politically motivated sophisticated offensive cyber operations can be constrained through the systematic adoption of security-by-design and vulnerability management practices by suppliers of digital products.

Specific hypothesis:
The feasibility of offensive cyber operations can be reduced by lowering the likelihood of real-world exploitation of vulnerabilities through the adoption of Product Security Incident Response Teams (PSIRTs) as a suppliers' good practice.

Individual hypothesis:

1. Leading suppliers of digital products have developed and implemented recognised good practices for secure design and vulnerability management throughout the product lifecycle.
2. Among these practices, the establishment and operation of a PSIRT contributes measurably to reducing the likelihood that disclosed vulnerabilities will be actively exploited in the wild as part of sophisticated offensive cyber operations.

1.5 Research Questions

This research is guided by the following questions:

- What security risks arise from vulnerabilities in digital products, particularly in relation to their exploitation by states and state-backed actors?
- How are these risks currently addressed through diplomatic initiatives, regulatory and legal frameworks, and industry-led practices?
- What secure development, vulnerability management, and coordination practices are employed by leading suppliers of digital products?
- Which of these practices are commonly regarded as effective in reducing vulnerabilities and limiting their exploitation?
- What methodologies and indicators – existing or proposed – are available for assessing the effectiveness of such practices?
- Which methodological approach is most appropriate for evaluating practices aimed specifically at reducing the likelihood of vulnerability exploitation in the wild?
- What insights does the application of such a methodology to real-world data provide regarding the effectiveness of a selected practice in mitigating active vulnerability exploitation?
- To what extent can the results of this assessment inform conclusions about the impact of suppliers' security practices on the feasibility and cost of conducting offensive cyber operations?

1.6 Scientific Contribution

The primary scientific contribution of this dissertation lies in the formulation and empirical testing of a novel methodological model for assessing the effectiveness of corporate security-by-design and vulnerability management practices in reducing the likelihood of vulnerability exploitation in the wild. The proposed model operationalises a set of vulnerability-related parameters to enable structured evaluation of how specific organisational practices affect real-world exploitation risk. As such, it provides a reusable analytical framework that can be applied to future practices, theoretical models, and policy interventions aimed at improving digital product security. The model also contributes conceptually by highlighting which parameters are both observable and analytically meaningful for assessing security outcomes, thereby opening new avenues for methodological refinement.

A central empirical contribution is the application of the proposed model to a major industry practice – PSIRTs – using publicly available datasets on disclosed and actively exploited vulnerabilities. This simulation-based analysis demonstrates that the presence of a PSIRT within a supplier organisation is associated with a statistically significant reduction in the likelihood of vulnerability exploitation in the wild. Beyond confirming the effectiveness of PSIRTs – while recognising the limits of causal inference inherent to observational designs – the study establishes a replicable approach to testing other security practices by adjusting datasets, parameters, or contextual factors.

To support the development and validation of the model, the dissertation further contributes by systematically collecting, classifying, and analysing security-enhancing practices employed by leading global technology companies, alongside academic and industry methodologies for assessing their effectiveness. It maps the global distribution of PSIRTs across industries, regions, and company sizes, and identifies key organisational, economic, and regulatory challenges associated with their establishment and continuous operation. This empirical mapping enables comparative analysis of best practices across multiple dimensions and grounds theoretical discussions in observable industry realities.

In addition, the dissertation advances methodological thinking by critically examining vulnerability-related parameters relevant to exploitation risk, such as severity scores, remediation timelines, exploit availability, and external factors. This analysis provides a foundation for future research that may explore alternative parameter sets or refine existing ones, potentially yielding models tailored to different types of practices or threat environments.

By linking the empirical findings of the proposed model to broader discussions on cyber armament and offensive cyber operations, the dissertation contributes an interdisciplinary perspective that bridges technical cybersecurity practice with international security studies. It demonstrates how improvements in digital product security can raise the cost, complexity, and uncertainty of conducting politically motivated offensive cyber operations, thereby contributing – albeit indirectly – to strategic stability in cyberspace. In practical terms, the model offers policymakers, regulators, and industry leaders an evidence-based tool to assess, incentivise, and prioritise security-by-design practices through both regulatory and market-driven mechanisms.

Finally, the dissertation contributes to theoretical development by integrating technological, organisational, political, and economic dimensions of cybersecurity. It connects corporate decision-making and engineering practices with state-level security dynamics, illustrating how micro-level organisational practices can have macro-level implications for cyber conflict and international relations. This approach encourages further cross-disciplinary research at the intersection of computer science, security studies, and political science.

1.7 Structure of the Thesis

The following chapter, *Chapter 2*, establishes the interdisciplinary theoretical and contextual foundations of the research by reviewing literature on cyber armament, OCO, and the exploitation of vulnerabilities in digital products. It also examines technological, economic, regulatory, and political approaches through which states and other actors seek to raise the costs of OCC and constrain their use.

Chapter 3 analyses and classifies existing practices for improving the security of digital products and reducing vulnerabilities. It reviews academic models, international standards,

regulatory frameworks, and industry-led SDL practices, with particular attention to vulnerability management and coordination mechanisms such as PSIRTs.

Chapter 4 examines existing industry and academic methodologies for assessing the effectiveness of security-by-design and vulnerability management practices. It identifies their limitations in evaluating real-world exploitation risks and formulates key technical, organisational, and contextual parameters relevant to assessing the impact of such practices on OCC.

Chapter 5 introduces a novel verification methodology for assessing the effectiveness of good practices in reducing vulnerability exploitation in the wild. The chapter presents the methodological design based on a matched case-control observational study and explains the selection of variables, datasets, and analytical techniques.

Chapter 6 operationalises the proposed methodology by applying sampling, matching, and bootstrapping techniques to analyse real-world vulnerability data and control for external influences. It introduces the selected risk-reduction metrics and statistical tests used to evaluate the robustness of the results, and to assess the role of key confounding factors.

Chapter 7 presents and analyses the results of the simulation, assessing their statistical reliability and practical significance. It evaluates the impact of PSIRT presence on vulnerability exploitation likelihood and examines the influence of confounding factors such as proof-of-concept availability and industry type.

Chapter 8 discusses the findings in a broader technical, organisational, and geopolitical context. It critically reflects on the role of industry practices in constraining OCO and cyber armament, identifies structural limitations, and explores implications for policy, regulation, and international security.

The concluding chapter, *Chapter 9*, summarises the research problem, proposed model, and key findings, and clearly articulates the scientific contributions of the study. It also offers practical recommendations for industry and policymakers and outlines directions for future academic research.

2 Review of the Research Field

Driven by pervasive digitalisation, nearly all segments of modern society – from private life and business to education, culture, and international relations – increasingly depend on information and communication technologies (ICT). Businesses rely on complex software systems for customer management, e-commerce platforms, and cloud environments, while industries rapidly integrate Internet of Things (IoT) devices combined with advanced telecommunications networks such as 5G. Critical infrastructures, including transportation, healthcare, and energy, increasingly depend on interconnected and 'smart' operational technologies (OT), where modern software solutions interface with sensors, control modules, and automation systems, often accessible remotely.

However, nearly all software – and much of the hardware – underpinning this digital ecosystem contains flaws that create vulnerabilities in both technical systems and society as a whole. This stems from a market dynamic that prioritises rapid innovation and product release over security, leaving vulnerability identification and patching to later stages. At the same time, the mass use of AI-assisted coding and unchecked integration of third-party open-source components propagate vulnerabilities across entire supply chains.

The exploitation of such vulnerabilities enables diverse threat actors to conduct cyberattacks. As societies become more dependent on digital systems and supply chains grow in complexity, malicious actors – including criminals, political groups, and states – have become increasingly capable and ambitious. States, in particular, with vast resources and strategic motivations, exploit or store critical vulnerabilities to conduct espionage or offensive cyber operations. While it is unlikely that states will abandon the use of vulnerabilities as part of their cyber arsenals, an important question arises: how can the cost and feasibility of cyber operations be altered, to result with decreased exploitation of vulnerabilities?

This chapter provides an overview of the existing literature from the field and theoretical foundations relevant to understanding the escalating cybersecurity risks of vulnerability exploitation, particularly by the state-related actors. It starts with analysing key assets, threats, and vulnerabilities. The chapter then reviews the political, diplomatic, and regulatory responses towards increasing the costs of cyber-armament through vulnerability exploitation – from international multilateral frameworks and norms, via multistakeholder platforms and principles, to national regulatory frameworks and trends – as well as the incentives and roles of the private sector in underpinning those. This foundation enables the subsequent chapters to focus more precisely on security-by-design approaches and good practices for strengthening the security of digital products, including industry models, international standards, and vulnerability management processes such as PSIRTs, as well as methodologies for evaluating their effectiveness.

2.1 Growing Cybersecurity Risks

As society continues to advance through digitalisation, it simultaneously faces threats that are becoming increasingly sophisticated and pervasive. Brunner and Sutter (2008) argue that cyberspace follows a pattern similar to other security domains, where security challenges are not direct, certain, and immediate. Instead, they should be perceived as risks: indirect, unexpected, and unplanned occurrences that manifest in the future, materialising only when they actually take place.

The digital realm fundamentally revolves around the flow and manipulation of digital data, which means that the risks of cyberattacks primarily reflect on the core components of information security: confidentiality, integrity, and availability of data (often referred to as the CIA triad) (Stamp, 2011). According to Ghazouani et al. (2014), leading literature employs mathematical formulas to assess information security risks (denoted as R), incorporating threats (T), vulnerabilities (V), and the value of assets that may be compromised (A) – and thus the related impacts of an attack – into a multiplicative model:

(1) $$R = T \times V \times A$$

This formula underscores the interdependence of these elements, indicating that the escalation of any single factor can lead to a significant increase in overall risk. A closer examination of these variables reveals that cybersecurity risks are inevitably increasing. The growing value of digital assets, combined with the evolving profile and tactics of cyber adversaries and persistent vulnerabilities within systems, creates a landscape where organisations must remain vigilant and adaptive.

In the context of international security, the ramifications of increasing cybersecurity risks extend beyond individual organisations, threatening the stability and peace of nations. Cyberattacks targeting critical infrastructure, such as energy grids, transportation systems, and financial institutions, can lead to widespread disruption and chaos, undermining public trust and national security. The World Economic Forum (2025) highlights that as cyberthreats continue to evolve, they pose significant risks to global stability, with potential impacts on economic systems and societal structures. The interconnectedness of global systems means that vulnerabilities in one region can have cascading effects, potentially destabilising entire economies and geopolitical relationships.

Moreover, state and non-state actors are leveraging cyber capabilities as tools for hybrid warfare, espionage, and coercion, blurring the lines between traditional military engagements and cyber operations (Libicki, 2007). This evolving landscape necessitates a comprehensive understanding of the risks posed by the interplay of assets, threats, and vulnerabilities, particularly as they relate to the preservation of international peace and security.

2.1.1 Assets

One of the variables in the cybersecurity risk equation is the value of assets that may be compromised or disrupted through cyberattacks. This value has been increasing in tandem with society's growing dependence on digital products and services. Already in its 2013 report on information and communications technology (ICT) developments in the context of international security, the United Nations Group of Governmental Experts (UN GGE) confirmed the escalating threats to individuals, businesses, critical infrastructure, and governmental institutions, alongside the increasing damage caused by cyber incidents, particularly through the use of tools like botnets (UN, 2013).

Critical infrastructure (CI) – including the electric grid, transportation, healthcare, and communication systems – is a cornerstone of national and international security (Wiśniewski et al., 2022). Its integrity is paramount: disruption of these systems can compromise national security, destabilise economies, and even result in loss of life (Zwarts et al., 2025). Because of both their instrumental and symbolic value, attacks on CI have become an integral part of modern strategies of disruption and destruction Coward (2009).

The consequences of cyberattacks on CI extend well beyond immediate technical damage. They can trigger cascading failures in essential services, thus threatening public safety (Palleti et al., 2021). The 2015 cyberattack on Ukraine's energy grid vividly illustrates these dangers: it caused power outages affecting more than 230,000 residents and exposed the vulnerabilities of interconnected infrastructures (Whitehead et al, 2017; NATO, 2016). Similarly, the 2020 SolarWinds breach demonstrated how supply chain vulnerabilities can compromise both government and private sector networks worldwide, highlighting the global scope of CI risks (FireEye, 2020). In recognition of these threats, NATO and the EU have elevated CI protection as a strategic priority, embedding resilience against cyberattacks into their collective defence and cybersecurity strategies (NATO, 2016; EU, 2020).

Attacks targeting cloud and hybrid infrastructure, including smart grid systems, have increased in frequency and sophistication (Gunduz and Das 2020). Cloud environments present unique security challenges due to their distributed nature, shared infrastructure and responsibility models: Ali et al (2015) outline challenges on the communication, architectural, contractual and legal levels. Hybrid infrastructure compounds these risks by creating complex interdependencies between different cloud and on-premises systems, with specific challenges including secure movement of data and workloads across data centres, shared technology issues, and vulnerabilities which cascade from one component across the entire architecture (Reza et al 2019). Since cloud systems are vulnerable to a particular type of targeted and sophisticated attacks, the potential for large-scale disruption makes them attractive targets for state-sponsored actors (Vance 2014).

Other critical assets at risk include democratic processes and national governance systems, such as elections. Election process can be under cyberthreat, as the example of the US elections in 2016 most notably showcase (Stedmon, 2020). Thuli et al (Mkhwanazi et al., 2025) outline

threats of election data breaches (including hacking election campaigns, altering voters' data or vote counts), phishing attacks, and election service disruptions; importantly, they observe the African context, signalling that these threats are not limited to developing countries only.

Importantly, cyberattacks against CI could cause medium-to-long term political instability: the EU cybersecurity agency (ENISA, 2019) warns of potential for provoking chaos and anarchy, sowing social divisions, and undermining trust in democracy. Matzkin et al. (2023), however, found out that cyberattacks on critical infrastructure don't necessarily undermine public trust in government, but could even heighten it by amplifying anger.

Additionally, human lives are an invaluable asset that can be jeopardised by cybersecurity risks, particularly when healthcare systems are compromised or when critical information is manipulated. The Cyberspace Solarium Commission (2020) report underscores that the vulnerabilities exposed during the COVID-19 period had significant implications for the security of healthcare assets, which are essential not only for individual well-being but also for the broader public health landscape. A tragic example of the potential consequences of cyberattacks on human lives is the case of a cyberattack on a hospital in Germany in September 2020, which resulted in the death of a patient after emergency services were diverted due to the hospital's systems being rendered inoperable (Computer Fraud & Security, 2020).

The risks to human lives in the context of conflict are exacerbated by cyberthreats, particularly for civilians and humanitarian organisations. Civilians in conflict zones are increasingly vulnerable to cyberattacks that can disrupt essential services, such as healthcare and communication, thereby compromising their safety and well-being. The UN Open-Ended Working Group (OEWG) 2025 report highlights the importance of safeguarding civilian infrastructure from cyberthreats during armed conflicts, advocating for the means to protect civilians and humanitarian efforts from cyber operations that could cause harm (UN, 2025). Furthermore, the International Committee of the Red Cross (ICRC) has documented the significant risks faced by humanitarian organisations in conflict-affected areas, including undermining of their operations, disrupted supply chains, and exposure of sensitive data, ultimately jeopardising the safety of both aid workers and the populations they serve (ICRC, 2020).

The role of artificial intelligence (AI) in these assets is increasingly significant. AI technologies are being integrated into various sectors and the supply chains, enhancing operational efficiencies and automating decision-making processes. However, this integration also introduces specific vulnerabilities, as adversaries can exploit weaknesses in AI systems to manipulate outcomes, disrupt services, or compromise data and resources (Squillace et al., 2024). The potential for AI to be weaponised further complicates the landscape, as the risks associated with compromised AI systems could have far-reaching consequences for international peace and security (Goddy et al., 2024).

2.1.2 Threats

Threats in the cybersecurity landscape are characterised by the actors involved, their tactics, techniques and procedures (TTPs), and the emerging technologies they leverage. Threat actors range from state-sponsored groups to organised criminal networks, each employing distinct methodologies to achieve their objectives (Rid and Buchanan, 2014). The TTPs employed by these actors have grown increasingly sophisticated, with advanced persistent threats (APTs) demonstrating systematic targeting of vulnerabilities (Hutchins et al., 2011). Emerging technologies such as AI and large language models (LLMs) are creating new vectors for cyberthreats while simultaneously enhancing defensive capabilities (Brundage et al. 2018, Jada and Mayayise, 2024). This multidimensional nature of cyberthreats requires comprehensive analysis to understand their full implications for international security.

2.1.2.1 States as Threat Actors

The 2025 consensus report from the Open-Ended Working Group (OEWG) and the 2015 report from the United Nations Group of Governmental Experts (UN GGE), highlight a paradigm shift in the nature of cyberthreats, particularly due to the increasing armament of states in cyberspace (UN, 2015; UN, 2025). This militarisation poses substantial risks to international stability and security, as it facilitates politically motivated complex OCO that can lead to significant geopolitical and geoeconomic consequences.

States are actively developing OCCs to complement the conventional warfare rather than substitute it (Gartzke, 2013). Research by the Digital Watch Observatory (Radunović and Gavrilović, 2025) systematically maps publicly available evidence – such as official policies, strategic documents, and authoritative statements – showing that a growing number of states have developed, or are actively developing, OCCs (Figure 1). The mapping also captures indications of political intent or commitment to OCC development based on credible media reporting and assessments by the technical community. In practice, however, the number of states possessing or pursuing OCC is likely higher than documented. The increasing availability of 'off-the-shelf' attack tools – both commercial products and those offered through illicit markets – has lowered the entry barrier, enabling even resource-constrained states to initiate offensive cyber programmes. Moreover, some states may deliberately refrain from publicly acknowledging such capabilities, while others may frame OCC-related investments under the pretext of cyber defence. Taken together, these dynamics suggest that cyber armament is not only widespread, but also expanding.

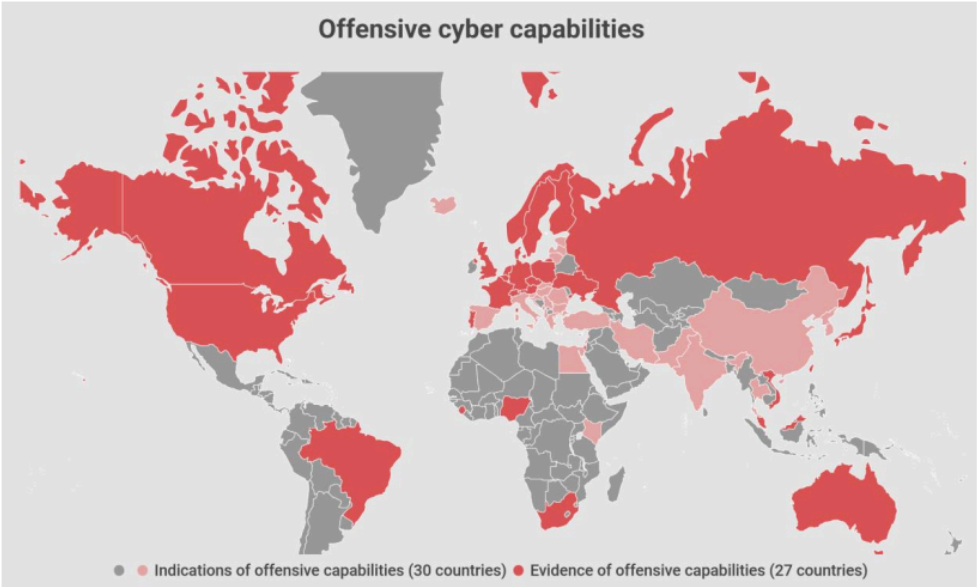


Figure 1: Map of states with offensive cyber capabilities (Radunović and Gavrilović, 2025)

Yet what makes OCCs more distinctive is their application in context of hybrid conflict scenarios, where cyber operations are combined with information campaigns and economic coercion (MSC, 2015; Lund, 2024). They can be deployed for subversion and disruption of social, economic, and political stability of target states while remaining short of open conflict (Maschmeyer, 2022), or as part of the intelligence contest, i.e. to conduct advanced espionage and achieve informational advantage (Baram 2023).

OCCs provide significant strategic value to states both in terms of contributing to power balance and conducting deterrence (Eriksson and Giacomello 2014). They allow states to target both military and civilian systems with limited or no casualties at relative low costs, while achieving economic, operational, and psychological effects (Smeets 2022). The real effectiveness of OCC in achieving such strategic advantages and deterrence, however, remains constantly questioned due to the fast-pace innovations in technology and tactics (Mitchell and Pytlak 2020).

Importantly, OCCs heavily rely on exploiting vulnerabilities in digital infrastructure. Mladenović and Radunović (2018) argue that vulnerabilities form the very foundation of OCC, as successful cyberattacks materialise only when the attacker's capabilities align with the opportunity to exploit existing weaknesses. Importantly, the concept of vulnerability in this context extends beyond purely technical flaws in system design and implementation to include human and organisational factors related to system management and use.

To benefit from OCC, states must operate in a domain of exploitation, which is characterised by persistent engagement and exploitation of vulnerabilities in digital systems of their

counterparties – for offense, defence, or deterrence (Dong et al., 2025; Fischerkeller et al., 2022). This, in turn, demands continuous weaponisation of vulnerabilities, which both undermines the economic model of digitalisation, which is based on trust, and contributes to the destabilisation of international security (Microsoft, 2017).

The threat landscape is further complicated by the role of state-sponsored or state-supported actors – companies, criminal, and hacker groups (Stoddart 2022). Research indicates that states frequently employ such groups as deniable proxies to conduct cyber operations while allowing them plausible deniability (Lin 2016). Further, the convergence between criminal and state-sponsored actors increases as criminal tools and infrastructure – those for ransomware, DDoS attacks, and data breaches – are repurposed for geopolitical objectives (Tomer, 2025, Lavorgna, 2023). The use of proxies creates a complex attribution environment, undermining accountability mechanisms and eroding trust in international norms (Rid and Buchanan, 2014; Egloff and Smeets 2021); the advances in AI techniques, however, promise to improve attribution accuracy (Prasad et al., 2025).

2.1.2.2 Tactics, Techniques, and Procedures

The evolution of tactics, techniques, and procedures (TTPs) among cyber adversaries is a central driver of the increasing threat landscape. The MITRE ATT&CK framework serves as a comprehensive reference, mapping and clarifying the diverse range of adversary TTPs (Strom et al., 2018). As adversaries adapt, the sophistication of cyberattacks continues to escalate, with state-sponsored and organised criminal groups refining their methods across a spectrum of attack types, including ransomware, wiper malware, zero-day exploits, and supply chain intrusions (Admass et al., 2024; UN, 2025; Sophos 2025; IBM 2025).

Ransomware attacks, in particular, have emerged as a prominent and persistent threat. The US Cybersecurity and Infrastructure Security Agency (CISA) highlighted a significant rise in ransomware incidents affecting critical infrastructure in 2021, underscoring the vulnerability of essential services to disruption by malicious actors (CISA, 2022). While some cybersecurity firms have reported a slight decline in the number of ransomware incidents in recent years, ransomware still constitutes nearly one third of all malware cases and increasingly targets crucial sectors such as manufacturing, healthcare, and finance (Sophos, 2025; IBM, 2025; TrendMicro, 2023).

The 2025 consensus report from the United Nations Open-Ended Working Group (OEWG) expresses particular concern about the increasing frequency, scale, and severity of ransomware attacks, specifically targeting CI, causing harm and disrupting essential services in sectors such as healthcare, finance, and energy (UN, 2025). Such attacks can have far-reaching consequences; for example, the 2022 ransomware attack on Costa Rica's national digital and data services prompted a state of emergency and led to market and governance failures, a phenomenon described by Simon and Haklai (2025) as 'Denial-of-Governance'.

Transitioning from ransomware, another area of growing concern is the proliferation of Cybercrime-as-a-Service (CaaS) models. These business structures allow threat actors to sell or lease malicious tools and services, lowering the technical barriers to entry for a wide array of adversaries. Ransomware-as-a-Service (RaaS) exemplifies this trend, enabling less skilled actors to launch sophisticated campaigns (Akinsuli, 2021). Double extortion, in which attackers not only encrypt data but also threaten to leak it, has further complicated the landscape. According to (Sophos, 2025), the proportion of organisations subjected to extortion doubled in 2025 compared to 2024, intensifying the risk of sensitive data leaks from public authorities and national security bodies, with direct implications for international security.

Closely linked to these developments is the emergence of supply chain attacks as a particularly insidious threat vector. By targeting software dependencies and third-party vendors, adversaries can compromise a broad array of downstream users (Hammi and Zeadally, 2023). The SolarWinds incident remains a defining example, where attackers exploited trusted update mechanisms to distribute malware to thousands of organisations simultaneously (Alkhadra et al., 2021). Such attacks exploit the inherent trust placed in software supply chains, necessitating enhanced vetting of third-party components and robust software bill of materials (SBOM) practices (Cyberspace Solarium Commission, 2020).

Zero-day attacks represent another formidable challenge, as threat actors exploit previously unknown vulnerabilities in widely used software before patches become available. The commercial market for zero-day exploits has expanded considerably (Google, 2023; Perlroth, 2021), with nation-states increasingly stockpiling such vulnerabilities for potential future conflicts (Lin, 2016). This trend creates significant windows of opportunity for attackers to compromise high-value targets and underscores the imperative for timely vulnerability disclosure and patch management. The exploitation of a zero-day vulnerabilities is of particular relevance for OCO, although it is a time-consuming approach (Sigholm and Larsson, 2014).

Artificial intelligence (AI) has also become a transformative factor in the evolution of TTPs, offering both defensive and offensive capabilities. On the one hand, AI-driven systems enable defenders to detect anomalies, automate processes, identify emerging threats, and manage vulnerabilities (Jada and Mayayise, 2024). On the other hand, adversaries are increasingly leveraging AI for both reconnaissance and attack automation (Brundage et al. 2018). Guembe at al. (2022) report that over half of AI-driven techniques are observed in the access and penetration phase of the cyber kill chain, but they also contribute significantly to exploitation, command-and-control, reconnaissance, and delivery of payload. Large language models (LLMs), such as ChatGPT, have been misused for fraud and impersonation, crafting convincing phishing or spear-phishing campaigns, assisting in malware generation, and supporting sophisticated social engineering schemes, often without requiring advanced technical expertise; moreover, criminal networks have developed ‘jailbreak-as-a-service’ to manipulate legitimate AI systems, bypass embedded safeguards, and create malicious outputs (Europol, 2023). These developments underscore that AI is not only amplifying the speed and reach of cyberattacks but also complicating attribution and defence, thereby reinforcing its emerging role as a dual-use enabler within the modern TTP landscape.

Advanced Persistent Threat (APT) groups exemplify the evolution and sophistication of threat actor tactics. These groups continue to develop and refine their TTPs. They primarily focus on data exfiltration and the disruption of critical systems (González-Manzano et al., 2022), employing sophisticated methods such as obfuscation, proxy routing, false flags, cross-jurisdictional infrastructures, and manipulation of geopolitical narratives – all of which make attribution increasingly difficult (Prasad et al., 2025). APTs typically operate with substantial resources and conduct long-term, targeted campaigns, often under the direction or support of state entities (González-Manzano et al., 2022). Their TTPs differ markedly from those of less sophisticated threat actors, demonstrating greater technical capability and operational discipline. Ahmad et al. (2019) further distinguish strategically motivated APTs (S-APTs), whose objectives align with the strategic interests of third parties, such as nation states or organised criminal networks. These actors exhibit high adaptability, continuously retooling and evolving their tool chains over time. Notably, Lemay et al. (2018) provide evidence that APT capabilities are no longer limited to a few dominant states but have proliferated among emerging actors, particularly in the domain of cyber espionage.

2.1.3 Vulnerabilities

The third variable that significantly influences the growing risks in information security is the presence of vulnerabilities within digital infrastructure, products, and services. Vulnerability exploitation and remains one of the most frequent technical root causes of cyber incidents. According to recent industry analyses, exploitation of already known vulnerabilities accounted for approximately 32% of documented attacks (Sophos, 2025), while four of the ten most frequently discussed vulnerabilities on dark web marketplaces are actively leveraged by sophisticated actors such as Advanced Persistent Threat (APT) groups (IBM, 2025).

Exploitation of vulnerabilities – often through the deployment of malware to gain unauthorised access to computer systems – constitutes one of the common, though not always indispensable, components of offensive cyber operations (Uren et al., 2018). Understanding how and why vulnerabilities emerge – and the factors that determine their likelihood of exploitation – is therefore essential for assessing and mitigating cyber risk.

2.1.3.1 Origin and Types of Vulnerabilities

In the business and technical context, the ISO/IEC 27000:2018 standard defines a vulnerability as a weakness of asset or control that can be exploited by one or more threats (ISO/IEC, 2018). Similarly, the U.S. National Initiative for Cybersecurity Careers and Studies (NICCS) defines a vulnerability as a weakness that exposes an asset or organisation to potential exploitation (NICCS, n.d.). In a broader, political interpretation, the U.S.–Russia bilateral glossary of cybersecurity terminology defines a cyber vulnerability (Rus. *Киберуязвимость*) as a characteristic of a cyber entity or object that is, or may become, susceptible to abuse or cyber operations (Godwin et al., 2014). These definitions indicate that vulnerabilities should

be understood within a broad conceptual framework that extends beyond purely technical weaknesses in software or hardware.

From a narrower technical perspective, a vulnerability may be defined as a flaw or defect that exposes software or operating systems to potential attacks or misuse (Pawlak, 2014). Defects can also arise in the design or implementation of hardware components. Indeed, nearly all digital products – software, hardware, or integrated systems – contain certain weaknesses introduced during their design, development, or deployment. Increasing system complexity, coupled with market pressures for rapid release cycles, leads many companies to knowingly deploy products containing hundreds of vulnerabilities, with the expectation that these will be patched later. Hathaway (2019) aptly refers to this as the 'field it fast, fix it later' ethos, which prioritises speed to market over product security.

Vulnerabilities can emerge at nearly every stage of a product's lifecycle: from design and planning, through coding or manufacturing, to deployment and integration. McGraw (2006) developed a taxonomy of common coding errors that can introduce vulnerabilities during software development. Building on this, Williams et al. (2018) classify vulnerabilities into two broad categories: implementation bugs and design flaws. Implementation bugs – such as buffer overflows or SQL injection – are often easier to detect and patch; moreover, they may remain undiscovered and thus not exploited for years. In contrast, design flaws – such as weaknesses in error handling or recovery mechanisms – are systemic and can only be remedied through redesign. Application-level vulnerabilities have consistently ranked among the top concerns of information security professionals, closely followed by configuration errors in networks and systems (Suby and Dickson, 2015).

A particular concern arises with vulnerabilities in open-source software (OSS). Compared to in-house or commercial off-the-shelf solutions, OSS offers clear advantages in terms of project metrics – time, cost, effort, and quality – as well as external factors such as market trends, availability of the source code, and licensing (Badampudi et al., 2016). However, these benefits are offset by growing security risks. The widespread use of OSS components and libraries across sectors means that vulnerabilities can propagate throughout the digital supply chain. High-profile cases such as Ripple20 and Amnesia:33 revealed numerous vulnerabilities in TCP/IP libraries used across Operational Technology (OT) and Internet of Things (IoT) products in industries ranging from healthcare and transportation to energy and retail (Kol and Oberman, 2020; Santos et al., 2020). These incidents underscore the systemic nature of software vulnerabilities and the challenges of securing the modern, interconnected ecosystem of digital products.

2.1.3.2 The Risk of Vulnerability Exploitation in the Wild

While identifying and classifying weaknesses is critical, the actual threat to organisations and infrastructures materialises through active exploitation in real-world cyberattacks ('in the wild'). The risk of such exploitation in operational contexts thus demands focused scrutiny.

To assess the risks associated with discovered vulnerabilities and to prioritise remediation efforts, ICT system operators and vendors rely on different models that estimate both the potential impact and the likelihood of exploitation. The Common Vulnerability Scoring System (CVSS) remains the most widely adopted open standard for qualifying the severity of vulnerabilities (FIRST, 2018). It measures both exploitability – the ease with which a vulnerability can be used by an attacker – and the impact that exploitation may have on the confidentiality, integrity, and availability of systems and data. However, despite its ubiquity, CVSS is often criticised for its limited predictive power.

Security defenders lack a common understanding of how CVSS scores correlate with the actual probability of exploitation in the wild (Allodi et al., 2018). In practice, a high CVSS score does not necessarily imply a high likelihood of exploitation: other contextual factors, such as the public availability of exploit code or the market demand for exploitation, are often more decisive (Allodi and Massacci, 2014). Furthermore, CVSS largely ignores real-time indicators of risk (Balsam et al., 2023), which limits its utility for operational decision-making. The newest version, CVSS 4.0, attempts to address some of these shortcomings by incorporating variables such as attack vector, environmental context, and probability of exploitation (FIRST, 2024), but its adoption remains limited.

In response to these limitations, several complementary models have been developed to help defenders prioritise their responses more effectively. The Exploit Prediction Scoring System (EPSS), for instance, applies machine-learning algorithms to historical exploitation data and intrusion detection alerts to predict whether a given vulnerability is likely to be exploited (Jacobs et al., 2023). Another framework, the Stakeholder-Specific Vulnerability Categorization (SSVC), tailors prioritisation decisions to the specific context of an organisation by evaluating parameters such as vulnerability exposure, active exploitation status, and the potential impact on public safety or critical operations (Spring et al., 2021). Despite their methodological sophistication, these models are still not widely implemented in industry practice – often due to their data requirements, computational complexity, and the inertia of established CVSS-based processes.

Importantly, traditional scoring models like CVSS treat vulnerabilities as static properties of software systems, neglecting the dynamic factors that shape their exploitation risk in practice. Empirical studies have shown that the probability of exploitation increases dramatically once a functional exploit becomes publicly available – whether through underground markets or legitimate proof-of-concept (POC) repositories (Allodi and Massacci, 2014). In fact, prioritisation policies that explicitly account for the presence of exploit code or POC evidence can reduce risk exposure by up to 80%. Further research confirms that the social and informational environment surrounding a vulnerability strongly influences its exploitation likelihood.

Vulnerability intelligence plays a pivotal role in reducing the likelihood that threat actors – including cybercriminals and state-supported groups – successfully exploit software and

hardware weaknesses. National CERTs/CSIRTs and internal PSIRTs take into account feeds on exploit availability, dark-web exchanges, proof-of-concept publication, and real-world attack telemetry to prioritise patching, issue advisories, and coordinate incident response across sectors. For example, by tracking when exploit code becomes publicly available or when a vulnerability begins to appear in active campaigns, these bodies can mount pre-emptive mitigations, restrict exposure windows and raise the cost for attackers.

Mentions of a vulnerability in exploit databases, exploitation proof-of-concept (POC) repositories, or dark web sites (Almukaynizi et al., 2017), as well as increased discussion across social media platforms such as X (formerly Twitter), can serve as early indicators of rising exploitation risk (Sabottke et al., 2015). To better reflect this reality, Jung et al. (2022) expanded the CVSS temporal metrics by adding an exploit code maturity attribute, while Bulut et al. (2022) proposed a complementary 'weaponised exploit' parameter. Nevertheless, these refinements remain non-standardised and depend heavily on subjective interpretation.

The number of publicly disclosed vulnerabilities continues to increase each year (NIST, n.d.). While public disclosure places pressure on suppliers to issue patches more rapidly (Arora et al., 2006), the time window between disclosure and active exploitation has been steadily narrowing: reports show an average time-to-exploit of 5 days, down notably from 32 days in 2021-22 and 63 days in 2018 (Charrier and Weiner, 2024; Charrier and Semrau, 2023). These compressed timelines highlight how quickly vulnerabilities become weaponised and why intelligence-driven prioritisation and coordination via CERT/PSIRT mechanisms are crucial for reducing exploitation risk.

It is important to note that the mere presence of a vulnerability in a digital product might not be a risk in itself. Many disclosed vulnerabilities don't introduce any significant risk, since they do not bring sufficient return to the attackers exploiting them. Actually, Mell, et al. (2024) estimated that 92% of the vulnerabilities are not being constantly exploited. Exploitation of more useful vulnerabilities by malicious actors, on the other hand, rely on sufficient and timely information about it, and a slow and uncoordinated response of a supplier and the defence community in providing and distributing a patch. Findings of the Carnegie Mellon CERT support this view, suggesting that security of a digital product should not be judged by number of discovered vulnerabilities, but rather by the suppliers' capability and maturity to respond to vulnerabilities (Householder et al. 2020).

The presence and timely deployment of a fix, therefore, significantly affect the likelihood of active exploitation – but not necessarily in a linear way. The speed and quality of patch development, and the coordination of its distribution and deployment, are critical determinants of residual risk (Radunović and Grätz-Hoffmann, 2020). Organisational culture thus plays a major role: empirical studies show that the median time to patch varies substantially across sectors and is often measured in weeks or months (Nappa et al., 2015). Industry research reports a median about 33 days, yet their Kaplan-Meier curve reveals that at that time over 50% of vulnerabilities still remain unresolved (Voicu and Fouque, 2025), as depicted in Figure 2. The presence of a dedicated PSIRT can help mitigate this by improving internal coordination,

accelerating vulnerability triage and patch development, and enabling coordinated disclosure with affected stakeholders.

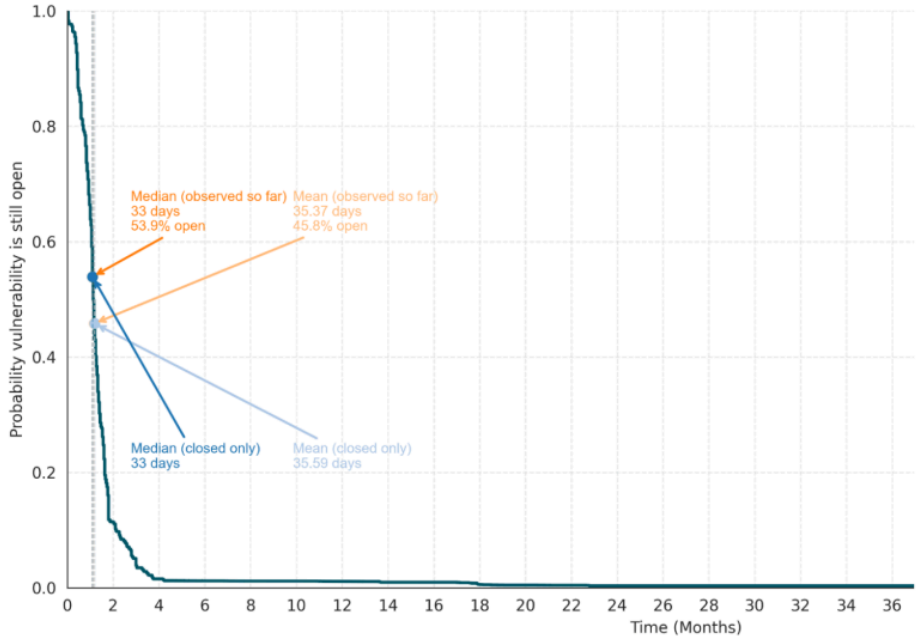


Figure 2: The Kaplan-Meier curve of the time-to-patch of vulnerabilities, versus a probability that a vulnerability is still open, also showing mean and median values. Source: Voicu and Fouque, 2025

Even when a fix is released, exploitation can continue. Arbaugh et al. (2000) applied their software vulnerability lifecycle model to three case studies – measuring exploit occurrences – and concluded that systems often remain vulnerable long after a patch has been released. Hathaway (2019) similarly notes that the 'gold standard' for patch deployment on the client side is approximately 30 days, yet many organisations – particularly in critical sectors – require significantly more time because they must test how patches affect essential operations. Besides, patch deployment depends on the timely distribution of information about vulnerabilities and fixes among those in risk, and their capacity to apply them. This extended window creates opportunities for attackers to exploit vulnerabilities immediately after disclosure. Combined with the rapid integration of insecure IoT systems and the growing cyber-armament efforts of states, Hathaway argues that companies must assume greater societal and legal responsibility to prevent vulnerabilities from emerging in the first place and shift from a patch-centric model toward ensuring higher levels of security before products reach the market.

Finally, the likelihood of exploitation is also impacted by the attacker's motivation and target selection strategy. Certain products, industries, and vendors attract disproportionate attention from threat actors due to their widespread use, strategic value, or perceived weaknesses. Most vulnerabilities exploited in the wild are found in operating systems, followed by Internet of

Things (IoT) and Internet of Medical Things (IoMT) devices, and then industrial control systems (ICS) and operational technologies (OT) (Baines, 2023). Predictably, the most frequently targeted vendors include Microsoft, Apple, Google, Oracle, and Adobe (Garrity, 2024). Vulnerabilities OSS represent another area of elevated risk, as they often suffer from limited long-term maintenance and can propagate downstream through software supply chains, multiplying their potential impact across dependent products.

The persistence and proliferation of vulnerabilities across digital products and infrastructures reveal a structural imbalance between technological innovation and security assurance. As exploitation becomes faster and more automated – driven by advanced threat actors, underground markets, and the availability of exploit kits – defensive actors often struggle to respond with comparable speed or coordination. Vulnerability intelligence, when efficiently shared through PSIRT, CERT, and CSIRT networks, can mitigate this gap by enabling faster prioritisation and patching, reducing the window of opportunity for malicious exploitation. However, the systemic nature of the problem means that purely technical countermeasures are insufficient. Reducing the likelihood and impact of exploitation ultimately requires a multilayered approach: strengthening national regulatory frameworks, fostering international cooperation and norm development, and embedding secure-by-design and coordinated vulnerability disclosure practices across the private sector. The next section therefore examines how these policy-oriented, diplomatic, and industry mechanisms can collectively raise the costs of offensive cyber operations and promote greater resilience of the global digital ecosystem.

2.2 Increasing the Costs of Cyber Armament and Vulnerability Exploitation

The ongoing 'cyber armament' of states, enhancement of OCCs, and the accumulation and deployment of software vulnerabilities – stems from a strategic imbalance: the perceived political and military benefits of cyberattacks often outweigh those of defence. As Libicki (2016) argues, the cyber domain remains offence-dominant because the cost of launching a sophisticated attack is relatively low compared to the resources required for effective cyber defence. In other words, each dollar invested in offensive capabilities can impose exponentially greater costs on defenders.

Dong et al. (2025) reflect on the emerging persistent engagement theory which claim that cyberspace is a specific strategic domain of exploration different from traditional domains of warfare, in which costs of state operations are low. Although high-end, escalatory cyber conflict has rarely happened, low-level, persistent operations are prominent (Lupovici, 2021). Such cyber operations allow states to carry out DDoS attacks, cyber espionage and data exfiltration, sabotage and other actions against another state, disrupting power dynamics without direct confrontation (Hunter et al., 2022; Maschmeyer, 2022). There, persistence directly depends on exploitation of vulnerabilities (Fischerkeller et al., 2022). Therefore, states' decisions to engage in cyber operations are shaped by perceived vulnerabilities of adversaries and the relative cost-benefit ratio of offence versus defence. Reducing the availability and exploitability of vulnerabilities therefore alters this strategic calculus.

To demotivate investment in cyber armament, this cost asymmetry must be reversed by raising the overall cost of conducting cyber operations. The final report of the U.S. Cyberspace Solarium Commission (2020) identifies cost imposition as one of three pillars of deterrence, recommending that adversaries be disincentivised through vulnerability reduction in critical sectors such as defence, nuclear, and energy infrastructures. Similarly, Morgus et al. (2017) argue that limiting the global supply of exploitable vulnerabilities can make it more costly and time-consuming for states to develop OCCs. Herr (2017) specifically suggests shortening the usability window of vulnerabilities, thereby forcing attackers to constantly retool and adapt their arsenals – a process that quickly becomes economically and operationally unsustainable. Findings by Ablon and Bogart (2017) support this reasoning: as in other markets, exploit prices rise dramatically when demand is high and supply is scarce, which naturally constrains access to advanced capabilities. These insights reinforce the rationale for both national and collective actions to reduce the incentives for cyber armament – including strengthening patch management, promoting responsible vulnerability disclosure, and developing cooperative frameworks for cyber stability.

In sum, reducing opportunities for vulnerability exploitation raises the operational, economic, and political costs of offensive cyber activity. Technical measures such as shortening exploit lifespans and improving vulnerability management directly constrain attacker capabilities. Economic mechanisms that limit exploit supply increase acquisition costs. And political and diplomatic initiatives – including the promotion of norms on responsible vulnerability disclosure and restraint in offensive operations – make cyber armament less attractive as a strategic tool. Collectively, these measures could discourage smaller states from developing offensive arsenals and compel larger powers to deploy their most valuable vulnerabilities more sparingly to avoid rapid detection and patching.

The subsequent sections consider how international diplomacy (UN processes and normative frameworks), national policies (vulnerability equities and regulation), and private-sector practices (coordinated disclosure, PSIRTs, secure-by-design and supply-chain governance) can jointly implement this cost-raising strategy and thereby reduce incentives for cyber-armament.

2.2.1 State-Level and International Efforts to Raise Costs

The previous section demonstrated that vulnerabilities remain a structural enabler of OCCs, fuelling a broader trend of cyber armament. The ambiguity in cyberspace lowers the political and operational costs of offensive action. In this context, international, national, and multistakeholder efforts to reduce systemic vulnerabilities serve an important strategic function: they directly raise the operational, economic, as well as reputational costs of conducting OCO and diminish the strategic utility of maintaining large exploit stockpiles.

2.2.1.1 The UN Framework

At the geopolitical level, cyber stability depends on establishing shared expectations of acceptable state behaviour (Finnemore and Hollis 2016; Henriksen 2019). Dunn Cavelty (2014) emphasises that, unless states adopt a strategic commitment to reducing systemic vulnerabilities, they will continue to pursue national security by expanding offensive cyber capabilities – ultimately degrading the security of global cyberspace. The UN framework on responsible state behaviour represents the most significant political mechanism attempting to reverse this logic.

The UN General Assembly's endorsement of the 2015 UN Group of Governmental Experts (GGE) report marked a decisive moment in the evolution of global cyber governance. For the first time, states collectively endorsed eleven non-binding norms of responsible state behaviour in cyberspace (Figure 3), including those directly addressing the handling, reporting, and mitigation of ICT vulnerabilities. Among these, norm 13(j) encourages states to promote responsible reporting of vulnerabilities and share information on remedies, while norm 13(i) calls for preventing the proliferation of malicious ICT tools and ensuring the integrity of ICT supply chains (UN, 2015). These norms established an initial conceptual link between state behaviour, vulnerability management, and the reduction of systemic cyber risks.

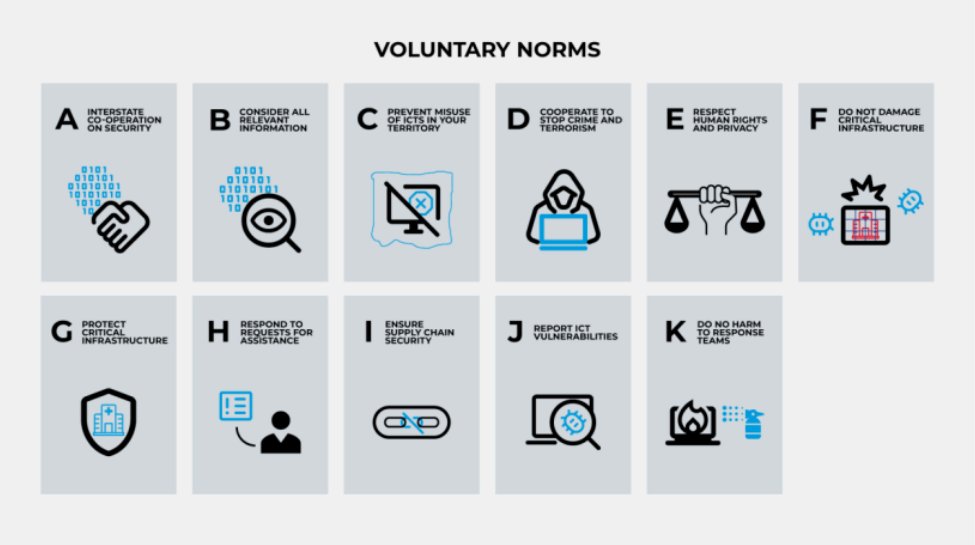


Figure 3. Eleven UN norms for responsible state behaviour in cyberspace adopted by the UN GGE in 2015 (Kazakova et al., 2025).

The significance of these norms has grown over time, in part because they emerged from a small group of states, led by the technologically advanced ones including the five permanent members of the United Nations Security Council: China, France, Russia, the United Kingdom, and the United States (Gavrilović et al, 2025). Cyber norms were later validated and extended

in a final GGE (UN, 2021), as well as through more inclusive multilateral negotiations: as the UN Open-Ended Working Group (OEWG) became the primary negotiating forum from 2019 onward, all UN member states – including many developing countries previously excluded from the GGE – endorsed the 2015 norm package. This broadened legitimacy was crucial: it transformed a specialised expert understanding of vulnerabilities and supply chain risks into a global political commitment, known as the UN Framework of responsible state behaviour in cyberspace, among other reinforcing the idea that reducing vulnerability exploitation is a shared responsibility.

The OEWG process also shifted emphasis toward implementation. Through its reports in 2021, 2023, and 2025 (UN, 2021; UN, 2025), the OEWG consistently reiterated the importance of supply chain integrity, vulnerability disclosure, and strengthened cooperation between national CERTs/CSIRTs (Ittelson et al, 2025).

The 2025 OEWG report (UN, 2025) goes further than previous ones by explicitly recognising the growing role of the private sector in managing vulnerabilities, acknowledging the risks associated with commercial spyware markets, and reaffirming that states should not knowingly allow their territory to be used for malicious cyber activity, including the deployment or sale of exploit tools. Importantly, this political document explicitly recognises (Section C, para 34(i)) security-by-design as an essential principle: states call on one another to ensure that ICT products are developed with security integrated from the outset, rather than relying solely on patching or reactive mitigation (UN, 2025). This marks a significant shift: vulnerability mitigation is now not just a matter of disclosure and repair, but of preventive design, thereby raising the cost of creating and maintaining exploitable weaknesses before deployment.

Confidence-building measures (CBMs) under the UN Framework have also increasingly incorporated elements related to vulnerabilities. The 2025 OEWG report recommends strengthening public-private partnerships (PPPs), mechanisms for vulnerability handling, recognising that private sector actors develop, detect, and disclose most vulnerabilities worldwide. The 2025 OEWG report promotes the establishment of a Global Points of Contact Directory (Section E, para 46(b)) and encourages cooperation among governments, industry and academia for threat-sharing and vulnerability disclosure (UN, 2025). This expands the norm landscape from state-centric obligations to a multistakeholder ecosystem, reflecting the reality that industry owns most infrastructure and generates much of the world's relevant technical expertise.

Although the consensus text of the final OEWG report represents significant progress, several deeply contentious issues raised during negotiations did not find full expression in the consensus language. During the negotiations, informal consultations and side-events, states and non-state experts voiced serious concerns about:

- the proliferation of commercial intrusion tools, exploit-as-a-service models, and hacking-as-a-service platforms – especially in the context of zero-day vulnerabilities

- the accelerating role of artificial intelligence in generating or automating exploits, training models on stolen data, producing spear-phishing campaigns at scale, and offering new vectors for vulnerability exploitation
- the absence of a global online platform (in contrast with national ones, like the US NVD platform) for sharing threat intelligence, vulnerability information and disclosure best-practices accessible to all states and companies, which could shorten exploit windows, democratise defensive capacity, and raise entry-costs for attackers
- the interaction of supply-chain risk, quantum-computing threats and the changing nature of 'legacy' vulnerabilities (Gavrilović et al, 2025)

These themes were omitted or diluted in the final consensus text due to geopolitical trade-offs and concerns about sensitive security capabilities. This underscores that OCCs – including the development, retention and operational use of vulnerabilities – remain strategically valuable for many governments, limiting the extent to which they are ready to formalise constraints. At the same time, the very fact that issues such as intrusion capabilities, zero-day stockpiling, AI-enabled exploitation and vulnerability markets were debated openly at the highest diplomatic levels reflects a growing recognition that such practices pose systemic risks. This presence of vulnerability-related concerns on the agenda of multilateral negotiations suggests that future rounds – particularly under the forthcoming permanent UN Global Mechanism – may return to these unresolved questions with greater urgency. Should that occur, these discussions could meaningfully shape international expectations around vulnerability handling and, ultimately, contribute to raising the cost of cyber operations and diminishing the strategic benefits associated with exploit arsenals.

2.2.1.2 Multilateral and regional processes

Regional organisations have complemented these UN processes. The confidence-building measures (CBMs) of the Organization for Security and Co-operation in Europe (OSCE) established structured mechanisms for exchange of information on cyber risks and cooperation among national incident response teams, including on disclosing vulnerabilities (Radunović, 2017). Similarly, the Association of Southeast Asian Nations (ASEAN) Regional Forum and the Organization of American States (OAS) adopted cyber CBMs that, while not always mentioning vulnerabilities explicitly, strengthen regional capacities for coordinated incident mitigation. These CBMs oriented to public-private partnership link normative expectations with operational practice: the more information flows between private firms and state incident-response agencies, the faster vulnerabilities are disclosed, the shorter the exploit window, and the higher the cost imposed on attackers.

The OECD Working Party on Security in the Digital Economy has taken a more operational and implementation-oriented approach, producing detailed guidance on how states and companies should manage vulnerabilities throughout the product lifecycle. Its recommendations on digital product security and coordinated vulnerability disclosure (CVD) (OECD, 2018; OECD, 2023) articulate concrete steps for embedding security-by-design,

improving supplier accountability, and institutionalising processes for receiving, assessing, and remediating reported vulnerabilities. By setting out governance models, minimum process requirements, and expectations for cross-sector cooperation, the OECD effectively translates high-level cyber norms into actionable policy frameworks. These outputs influence national regulatory reforms, inform industry standards, and contribute to emerging global convergence on best practices for reducing systemic vulnerability exposure.

Export-control regimes have also attempted to raise barriers to exploit proliferation. The Wassenaar Arrangement introduced controls on intrusion software and related exploit components, though the effectiveness of this measure remains contested. Herr (2016), however, warns that such controls may inadvertently hinder defensive research more than offensive capabilities due to the dual-use nature of cybersecurity tools.

2.2.1.3 Multistakeholder principles

Recognising that governments alone cannot secure the digital ecosystem, private-sector and civil-society actors have joined and undertaken substantial efforts to reduce vulnerabilities and strengthen product security. These initiatives translate political commitments into concrete supply-chain and lifecycle practices.

The Paris Call for Trust and Security in Cyberspace – initiated by France and Microsoft in 2018 and signed by more than 70 states and hundreds of companies – explicitly commits actors to strengthening the security of digital products and supply chains throughout their lifecycle, thereby reducing systemic vulnerabilities at scale. Its fifth principle calls for developing "ways to prevent the proliferation of malicious software and practices intended to cause harm", while the sixth aims to "strengthen the security of digital processes, products and services, throughout their lifecycle and supply chain." (Paris Call, 2018).

Similarly, the Cybersecurity Tech Accord binds more than one hundred global companies to prioritise security in the design, development, and deployment of their products, reflecting an industry-wide shift towards vulnerability-aware product governance (Cybersecurity Tech Accord, 2018). The Charter of Trust (2018), initiated by Siemens, IBM and others, goes further by promoting security-by-default configurations, transparent security assurances, and independent certification of digital components. These mechanisms, widely adopted by signatory companies, contribute to reducing latent vulnerabilities and increasing the effort required for attackers to weaponize exploit chains (Charter of Trust, 2018).

A multistakeholder process particularly relevant for addressing vulnerabilities and the supply chain is the Geneva Dialogue on Responsible Behaviour in Cyberspace, set-up by Switzerland in 2018. Unlike traditional intergovernmental forums, the Geneva Dialogue focuses on operationalising UN norms through non-state actors, especially with regard to vulnerability management, coordinated disclosure, supply-chain security, and secure-by-design practices.

The Geneva Manual (Kazakova et al., 2023) provides actionable guidance for governments and companies in implementing these norms.

More recently, the Pall Mall Process – a multilateral initiative launched by the UK – has sought greater oversight of the commercial spyware and intrusion industry, attempting to raise acquisition and operational costs for malicious actors and constrain the global market for exploit-based surveillance tools. However, Herpig and Paulus (2024) argue that the initiative overlooks a significant issue: the prominent role that governments themselves play in driving and sustaining the expansive, and often unregulated, surveillance market. This omission has led to concerns that the process may fail to address the root causes of the industry's unchecked growth and its impact on global cybersecurity and human rights.

2.2.1.4 National regulatory frameworks

At the national level, states increasingly recognise that vulnerability management and product security must be supported not only by voluntary practices but also by regulatory frameworks and institutional processes. Supply chain security, security of digital products and CVD has become in focus of the national cybersecurity authorities. In the United States, Executive Order 14028 (White House, 2021) introduced systemic supply-chain security requirements, mandated the use of Software Bills of Materials (SBOMs), and enhanced federal standards for secure development and vulnerability handling – significantly raising the cost of maintaining exploitability across ICT ecosystems. The EU NIS2 Directive (EU, 2022) placed responsibility on private sector entities to manage supply chain cybersecurity risks and to report security incidents, including through vulnerability disclosure processes, while the EU Cyber Resilience Act (EU, 2024) introduces further obligations for secure development, lifecycle vulnerability management, and vulnerability reporting for a wide range of digital products. Complementing these, labelling schemes for security of digital products introduced by Singapore, Germany, Finland, Republic of Korea and UK, for instance (CSAS, 2025; UK, 2019), raise baseline product security and limit the attack surface for exploit-driven operations.

Increasingly, national CERTs and authorities serve as voluntary or de-jure intermediaries between researchers and vendors, operating national CVD platforms, defining legal protections for security researchers, and establishing reporting obligations for essential-service operators (Kamara and Boom, 2022). NIS2 Directive also formally designates (one of) the national CERTs in each Member State to act as the coordinator of CVD and allows them to disclose vulnerabilities to ENISA's database (EU, 2022).

At the same time, states face strategic dilemmas when managing vulnerabilities internally. The US Vulnerabilities Equities Process (VEP) (White House, 2017) outlines how federal agencies decide whether to disclose or retain vulnerabilities discovered in commercial and other systems. While most vulnerabilities are disclosed for patching, a subset is retained for intelligence, law-enforcement, or military purposes. UK has introduced its own Equities Process in 2018 (UK, 2018b). This leads to a dual-use dilemma: when states retain

vulnerabilities, they exacerbate systemic ambiguity and inadvertently lower global security, allowing offensive incentives to proliferate. Increasing national involvement in CVD thus carries both stabilising and destabilising implications, depending on how governments balance disclosure and retention.

While Egloff and Smeets (2021) call for transparency on how states decide within their equities processes, and Romanosky (2019) even suggests a possible repeatable scoring system for VEP, in reality we haven't seen many improvements towards greater transparency and accountability.

Taken together, these international and regional frameworks, national regulations, and multistakeholder processes operate as a layered ecosystem of cost-imposition mechanisms. Although their implementation remains uneven and many frameworks still contain significant gaps, they nonetheless work toward reducing the supply of high-value vulnerabilities, shortening the shelf-life of exploits, hardening supply chains, and increasing detection and mitigation speed through strengthened CERT/CSIRT collaboration. They also elevate the political and reputational risks associated with stockpiling or misusing vulnerabilities.

Importantly, software and hardware development is no longer confined to the ICT sector alone: an increasing number of critical industries – such as finance, automotive, and energy – as well as local administrations and civil-society organisations now integrate or even produce and deploy digital products, often without the security maturity traditionally associated with the ICT domain. Since all such products can become targets of cyberattacks, including offensive cyber operations, the adoption of good practices for secure product development and deployment and the reduction of systemic vulnerabilities becomes essential for raising the overall cost of attacks and constraining cyber-armament dynamics more broadly.

In doing so, these initiatives collectively challenge the strategic calculus that currently underpins cyber armament. By making offensive cyber operations more expensive, less scalable, and less reliable, they gradually shift incentives toward defence, resilience, and responsible behaviour. Yet the effectiveness of these political and regulatory efforts ultimately depends on their translation into practical security measures within the products themselves, throughout their entire lifecycle and supply chain. As software and hardware development increasingly occurs across diverse industries beyond the ICT sector, the responsibility for implementing secure-by-design principles, coordinated vulnerability disclosure, and other internationally endorsed practices rests heavily on manufacturers and service providers. The next section therefore turns to the role of businesses in operationalising these norms and in embedding robust security practices into the design, development, and maintenance of digital products – an essential step for meaningfully reducing vulnerabilities and raising the overall costs of cyber exploitation.

2.2.2 Role of the Industry

As the primary producers, operators, and maintainers of the global digital ecosystem, businesses play an essential role in shaping the cybersecurity. Their decisions – ranging from product design and development to vulnerability disclosure and supply-chain management – directly influence the availability of exploitable weaknesses and thus the operational environment for OCCs. Understanding how industry actors engage with and implement international norms and principles is therefore a critical component of any broader effort to raise the costs of cyber armament and reduce the strategic value of vulnerability exploitation.

2.2.2.1 Industry as a Global Actor in Cyber Stability

Industry therefore occupies a particularly consequential position not only because it owns and operates most of the digital infrastructure, but also because it acts as the primary driver of technological innovation and the 'pace-setter' of the digital ecosystem (Kaufmann 2016). Since cyber operations transcend borders, the global reach of ICT and adjacent industries gives them an implementation capacity that states alone cannot match. When leading industry adopt consistent practices such as coordinated vulnerability disclosure (CVD) or secure-by-design development, they effectively raise the baseline of security worldwide, rather than merely within individual jurisdictions.

This global reach is particularly relevant for the implementation of cyber norms. Because cyberattacks typically originate across multiple jurisdictions and are executed remotely, the security posture of widely deployed commercial products becomes a de facto baseline for global cybersecurity. If major ICT vendors adopt robust practices – such as secure-by-design principles or CVD – these practices diffuse across the international market, effectively institutionalising norm implementation beyond the confines of national regulatory frameworks.

Industry's incentives to support such measures are not purely normative. Over the past two decades, companies have faced growing economic costs associated with cyberattacks, prompting an internal shift from treating cybersecurity as a technical add-on to viewing it as integral to business continuity, customer trust, and brand reputation (Anderson, 2001). These economic and strategic considerations converge with broader frameworks of corporate social responsibility (CSR), which emphasise how firms respond to both internal motivations and external pressures to act responsibly in areas where traditional regulation struggles to keep pace (Airike et al., 2016).

Taken together, industry's innovation capacity, infrastructure control, and global market presence position it as a central actor in the reduction of vulnerability-related risks. By shaping security architectures, implementing disclosure processes, and supporting the diffusion of norms, businesses directly influence both the supply of exploitable vulnerabilities and the feasibility of OCC. This makes them indispensable partners in the broader ecosystem of efforts

aimed at raising the costs of cyber armament and strengthening the resilience of the global digital environment.

2.2.2.2 Evolving Responsibilities in Norm Implementation

The growing scale and sophistication of cyberthreats brings awareness on how insecure products create systemic risks. This has led to a notable shift in corporate behaviour. Companies are increasingly expected not only to respond to incidents but to proactively prevent vulnerabilities, embed security across the product lifecycle, and contribute to the implementation of international cyber norms through their design decisions, disclosure practices, and supply-chain governance.

Existing literature identifies several concrete corporate roles in the implementation of cyber norms. Fairbank (2019) highlights its role in supporting attribution processes; meanwhile, the Geneva Dialogue mapping (Rizmal and Radunović, 2019) outlines responsibilities that directly relate to the vulnerability lifecycle: sharing information on vulnerabilities and best practices, contributing to standardisation efforts that embed security by design, improving end-user protection through greater privacy and integrity safeguards, and increasing transparency regarding product security and breaches. Discussions with industry leaders further emphasise commitments to strengthening critical infrastructure protection and more rigorous product testing (Eggenschwiler, 2018).

However, it is clear that voluntary and CSR efforts alone are insufficient, particularly given the scale and systemic nature of contemporary cyber risks (Maxwell and Barnsby, 2019). Scholars have long warned that manufacturers often lack transparency in the security of their products. Matwyshyn (2010) cautioned that software producers lacked sufficient transparency regarding product security and advocated a tripartite obligation: securing their code, warning promptly when vulnerabilities or exploitations arise, and providing effective patches. Hathaway and Savage (2012) went further, suggesting that governments should assign regulated liability to companies – clarifying processes for vulnerability disclosure, mandating early warnings, and ensuring producers assume responsibility when flaws create systemic risk.

2.2.2.3 Pressures Driving Security-by-Design

Research indicates that the cost and difficulty of patching vulnerabilities after deployment is rising (Dougherty et al., 2009), while systemic weaknesses in software architectures remain pervasive (Santos et al., 2017). As a result, companies are increasingly shifting from reactive patching toward embedding security earlier in the design and development process. This reflects the broader movement toward secure-by-design principles, lifecycle security, and resilience-focused engineering.

Viewed together, these regulatory, economic, and normative pressures align around a common expectation: that industry must take a more proactive role in reducing vulnerabilities and

operationalising international norms. By embracing secure development practices, improving transparency, and institutionalising vulnerability disclosure processes, businesses not only contribute to global cyber stability but also directly reinforce the broader cost-imposition strategies pursued by states. In doing so, they reinforce the broader cost-imposition strategies pursued by states, completing the link between international norms, national regulation, and operational measures to reduce the strategic utility of offensive cyber operations.

This naturally leads to the next chapter, which examines in greater depth how businesses operationalise these expectations through concrete good practices for digital products – particularly security-by-design, secure software lifecycles, and vulnerability management.

3 Industry Models and Practices for Reducing the Vulnerabilities in Digital Products

Efforts to reduce vulnerabilities and secure the broader technology supply chain involve a diverse set of actors – regulators, national authorities, OSS communities, cybersecurity researchers, civil society organisations, and academic institutions. Yet suppliers of digital products carry a particularly central responsibility: they design, develop, integrate, and maintain the hardware and software on which all other stakeholders depend (Kazakova et al., 2023). As (ENISA, 2023) notes, robust supplier-driven practices aimed at preventing and reducing vulnerability exploitation are indispensable for mitigating cyber risks and strengthening global cyber resilience.

Over the past two decades, the rising costs of patching, the persistence of architectural weaknesses, and the increasing weaponisation of software vulnerabilities have all reaffirmed the need to address security much earlier in the development lifecycle. This 'shift left' approach has driven the emergence of security-by-design as an economically viable systemic response (Waidner et al., 2014): a methodological framework that integrates security requirements, secure architectures, threat modelling, and verification mechanisms throughout all phases of product development. Rather than treating vulnerabilities as isolated defects to be fixed only post-deployment, this approach repositions security as a foundational engineering concern, complementing – rather than replacing – the ongoing need to detect and remediate vulnerabilities in deployed systems.

Against this backdrop, this chapter examines the principal conceptual and operational frameworks that shape contemporary approaches to securing digital products. It reviews the main academic models that introduced structured methods for analysing threats and designing secure architectures; surveys the international standards and regulatory instruments that codify secure development and vulnerability-handling practices; analyses the secure development lifecycle (SDL) practices adopted by leading technology suppliers; and outlines the organisational processes associated with vulnerability management and coordinated disclosure, in particular the role of product security teams (PSIRT). Together, these elements provide the groundwork for the subsequent assessment of how good practices can be evaluated and how effectively they reduce vulnerability exploitation in the wild.

3.1 Models and Methods from the Academic Sector

A substantial body of academic literature examines potential security models for product design and for specific system components. McDermott and Fox (1999), and later Sindre and Opdahl (2005), extended the classical use-case methodology for requirements engineering by introducing misuse (or abuse) cases alongside desired use cases, ensuring that security risks are captured directly within system requirements. Mouratidis et al. (2005) adapted the *Tropos* software-development methodology so that security becomes an integral element across the requirements, analysis, design, and implementation phases. Their approach introduces the

concepts of security constraints, entities, and dependencies early in requirements analysis as well as during architectural and detailed design.

Dougherty et al. (2009) compiled a collection of security patterns drawn from industry practice – templates describing generalised, reusable solutions to known security problems that can be applied across contexts to prevent vulnerabilities in code or minimise their impact. They categorise these patterns by abstraction level: architectural (e.g. privilege separation), design-level (e.g. secure chain of responsibility), and implementation-level (e.g. secure logging or removal of sensitive data). In a systematic mapping of secure software engineering (SSE) methods, Khan et al. (2020) analysed 116 primary studies and concluded that most existing models do not explicitly incorporate security techniques throughout the entire software lifecycle.

Numerous additional methodologies exist for individual components of SSE, or secure development lifecycles (SDL) – as a structured process for embed security into the very design of the software, rather than at the end. Shevchenko (2018), drawing on work at Carnegie Mellon University, mapped twelve threat-modelling methods used during product and service design to profile potential attackers, their objectives, and their methods, thereby generating structured catalogues of possible threats. Companies participating in the Charter of Trust for a Secure Digital World have, under Principle 3, formulated nineteen baseline requirements for security-by-default, including secure credentials, secure access mechanisms, protected data storage, secure update processes, mandatory password change upon first login, and masked password entry (Charter of Trust, 2019).

Santos et al. (2017) created the CAWE (Common Architectural Weakness Enumeration) catalogue, containing 224 common software architectural weaknesses, organised around established security tactics, and proposed techniques for their mitigation. Bunke and Sohr (2011) developed a tool dubbed *Bauhaus* which uses reverse engineering to reconstruct software architecture, detect security patterns, and propose methods by which developers can preserve these patterns through maintenance and updates.

Recent analyses of SSE and SDL also highlight organisational enablers such as explicit security metrics and strategies, compliance mechanisms, and ongoing security training (Mothanna et al., 2024). Agile development methodologies, which emphasise iteration, continuous learning, and rapid integration of feedback, have likewise been shown to improve software security when combined with structured secure-development controls (Patterson et al., 2023). Wee et al. (2024) report that organisations prefer flexible custom-made frameworks rather than SDL frameworks, as those better fit their internal development processes.

Taken together, academic models and methods provide a rich conceptual foundation for understanding how security can be embedded throughout the SDL. Yet, despite their analytical sophistication, these approaches often remain fragmented, highly specialised, or insufficiently aligned with the constraints and workflows of commercial development environments. As a result, they have had limited wider uptake in industry practice. This gap has prompted standards

bodies to translate early industry practices and academic insights into more operational, harmonised, and widely adoptable frameworks. The following section therefore examines how international and national standards codify security-by-design principles, establish common terminology, and provide structured processes for managing vulnerabilities and reducing systemic cyber risk.

3.2 International and National Standards and Recommendations

International technical and regulatory standards offer structured guidance for embedding security into digital products and services throughout their lifecycle. Within the International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC) ecosystem, several standards issued by the Joint Technical Committee ISO/IEC JTC 1/SC 27 set security requirements for manufacturers: ISO/IEC 27034-7:2018 provides engineering security requirements, testing and evaluation guidelines; ISO/IEC 29147:2014/2018 and ISO/IEC 30111:2019 establish internationally recognised processes for vulnerability disclosure and handling; and the 2022 revision of ISO/IEC 27001 introduced Annex A controls 8.27 (Secure Systems Architecture and Engineering Principles) and 8.28 (Secure Coding), requiring that security be embedded systematically across architecture, design, and implementation layers.

In parallel, the International Telecommunication Union (ITU) has developed a complementary body of standards, not directly addressing yet very relevant to security-by-design and vulnerability coordination. The ITU-T X.1500 series (Cybersecurity Information Exchange Framework) outlines mechanisms for structured sharing of security information, including vulnerability data. ITU-T X.1051/ISO 27011 provides sector-specific security management guidance for telecom environments, and ITU-T X.1205 offers a comprehensive overview of cybersecurity threats, vulnerabilities, and mitigation practices. Collectively, these ITU standards reinforce principles of SDL management, coordinated vulnerability disclosure, and cross-stakeholder information exchange, thereby complementing ISO and regional regulatory frameworks.

The Internet Engineering Task Force (IETF) was the first to set vulnerability disclosure standards through its *Responsible Vulnerability Disclosure Process* memorandum, which remains influential in shaping co-operation between researchers, vendors, and intermediaries (Christey and Wysopal, 2002). The joint International Society of Automation (ISA) and IEC standard serries ISA/IEC 62443, on the other hand, addresses standards in the industrial domain: it defines SDL requirements for digital products used in operational technology (OT) and industrial control systems. These include secure design principles, coding and deployment guidelines, verification and validation requirements, patch and bug management, and end-of-life considerations.

Beyond international standard-setting organisations, national and regional standard-setting bodies have developed frameworks that increasingly influence global practice. The US

National Institute of Standards and Technology (NIST) – a globally recognised authority in cybersecurity guidance – has developed the *Secure Software Development Framework* (SSDF), which provides terminology, organisational processes, design considerations, and guidance on responding to vulnerabilities and incidents (Dodson et al., 2020). NIST has also published security-by-design recommendations for Internet of Things (IoT) manufacturers (Fagan et al. 2020a; Fagan et al. 2020b).

The EU Agency for Cybersecurity (ENISA) has issued guidelines for SDL and good practices for IoT security, covering design, development, testing, deployment, and maintenance, while emphasising the role of technology, human factors, and organisational processes (ENISA, 2019). At the regulatory level, the *EU Cyber Resilience Act* (EU, 2024) establishes legally binding security-by-design and vulnerability management obligations for all digital products placed on the EU market, further operationalising these standards.

Singapore's Cyber Security Agency (CSAS), under its national Cybersecurity Act, has published additional SDL guidance, including a security-by-design framework demonstrating how conventional development cycles can be mapped onto secure development processes, accompanied by practical checklists for manufacturers (CSAS, 2017; updated 2022). The UK National Cyber Security Centre (NCSC) has defined eight core principles for improving and assessing the security of digital products. These include competence development, clean and maintainable code, secure development environments, continuous testing, and planning for vulnerability discovery and management, each accompanied by concrete implementation steps (UK, 2018a).

Emerging national certification and labelling schemes for security of digital products also introduce baseline security requirements and expectations for the industry. Notable examples are the ENISA's IoT Security Recommendations[4] from 2019 and Singapore's Cybersecurity Labelling Scheme (CLS)[5], but also Japan's JC-STAR IoT Security Label[6], or Finland's Cybersecurity Label[7] (now discontinued due to new EU-wide rules). These schemes sought to make product security visible to consumers and to incentivise manufacturers to adopt secure-by-design principles. Countries are also negotiating mutual recognition agreements, reflecting a growing recognition that fragmented national labels cannot support global supply chains.

Table 1 provides an overview of key international and national standards and frameworks for digital product security, including vulnerability disclosure and handling.

[4] Available at: https://www.enisa.europa.eu/sites/default/files/publications/WP2019%20-%20O.1.1.1%20Good%20practices%20for%20security%20of%20IoT.pdf
[5] Available at: https://www.csa.gov.sg/our-programmes/certification-and-labelling-schemes/cybersecurity-labelling-scheme/about/
[6] https://www.ipa.go.jp/en/security/jc-star/index.html
[7] Available at: https://tietoturvamerkki.fi/en/cybersecurity-label

	Organisation / Standard	Category
ISO/IEC	ISO/IEC 27034-7:2018	Application Security Standard
	ISO/IEC 29147:2014/2018	Vulnerability Disclosure Standard
	ISO/IEC 30111:2019	Vulnerability Handling Standard
	ISO/IEC 27001:2022 Annex A 8.27	Secure Architecture Requirement
	ISO/IEC 27001:2022 Annex A 8.28	Secure Coding Requirement
IEC	ISA/IEC 62443	Industrial/OT Security Standard
ITU	ITU-T X.1500 series	Cybersecurity Information Exchange
	ITU-T X.1205	Cybersecurity Overview Framework
	ITU-T X.1051 / ISO 27011	Telecom Security Management
IETF	IETF RFC - Responsible Vulnerability Disclosure	Disclosure Policy Guideline
National	NIST SSDF (SP 800-218)	Secure Development Framework
	NIST IoT Security Guidance	IoT Security Standard
	ENISA IoT Security Recommendations	IoT Security Guidance
	Singapore CSA – Secure-by-Design Framework	Secure-by-Design Framework
	UK NCSC Secure Development Principles	Product Security Principles

Table 1. Key International Standards and Frameworks for Digital Product Security

These international standards, regulatory frameworks, and technical guidelines form an increasingly coherent architecture for embedding security throughout the lifecycle of digital products. In practice, leading suppliers have already internalised many of these requirements, with industry practice both aligning with and actively informing the evolution of standards through feedback loops from PSIRTs, secure-development teams, and sector-specific communities. This mutual reinforcement between standardisation and practice provides a foundation for examining how the industry's own SDL models operationalise these principles in concrete organisational processes.

3.3 Secure Development Lifecycle Practices of Leading Suppliers

Leading global technology companies have developed SDL frameworks to systematise the ways to embed cybersecurity within every stage of product development. These models provide structured approaches for reducing the likelihood and severity of vulnerabilities, ensuring regulatory and standards alignment, and embedding security-by-design principles across complex engineering organisations. Several representative SDL models illustrate the diversity and maturity of current industry practice:

- Microsoft Security Development Lifecycle[8] remains one of the foundational industry frameworks. It prescribes a set of engineering practices aimed at reducing both the number and impact of vulnerabilities while controlling development costs. Introduced in 2004 already, the model embeds security tasks across all stages of the product

[8] Available at: https://www.microsoft.com/en-us/securityengineering/sdl/practices

lifecycle – including project initiation, definition of secure design requirements, threat and risk assessment, documentation, secure coding and testing policies, secure deployment, incident-response planning, and post-release security activities (Lipner, 2004; Howard and Lipner, 2006).

- Cisco Secure Development Lifecycle[9] offers a similar structure, combining technical tools, governance processes, and staff training across planning, development, validation, deployment, and operational phases. Introduced in 2013, its core components include defining security requirements, managing risks arising from third-party components, secure design and coding practices, security analysis, and vulnerability testing (Cisco, 2016).
- Adobe Secure Product Lifecycle (SPLC)[10], inspired by Microsoft SDL, is based on eight key stages: staff training and certification; requirements definition; secure design; development and testing; staging and stabilisation; deployment; operations and monitoring; and response to misuse, fraud, or security incidents. The SPLC applies to both software applications and cloud services (Adobe, 2018).
- Juniper Networks Secure Development Lifecycle[11] follows closely the Microsoft SDL structure, incorporating secure-coding education, secure architectural design, threat modelling, penetration testing, security reviews, and structured incident-response planning (Juniper, 2015).
- Software Assurance Forum for Excellence in Code (SAFECode)[12] aggregates core secure-software development practices, including secure design, secure coding, TPC risk management, testing and validation, vulnerability disclosure, and incident response (SAFECode, 2018).
- Open Web Application Security Project Software Assurance Maturity Model[13] (OWASP SAMM) provides a widely adopted methodology for assessing and improving an organisation's software-security maturity. The Software Assurance Lifecycle in the SAMM 2.0 framework outlines practices across governance, design, implementation, verification, and operations and incident response (OWASP, 2019).
- The Industrial Internet Consortium[14] (IIC) offers best practices for software trustworthiness across the full development lifecycle, placing particular emphasis on deployment, configuration, use and operational reliability in industrial environments.

[9] Available at: https://www.cisco.com/c/en/us/about/trust-center/technology-built-in-security.html
[10] Available at: https://www.adobe.com/trust/security/product-security.html#
[11] Available at: https://www.juniper.net/us/en/security/sdl/#gsc.tab=0
[12] SAFECode is an industry alliance whose members include leading technology companies such as Siemens, Dell, Microsoft, Google and Intel. The organisation focuses on identifying and promoting practices for the development of more secure and reliable software, hardware and services. More information is available at: http://www.safecode.org/index.php
[13] OWASP is a non-profit foundation built around knowledge-sharing within a community of tens of thousands of cybersecurity experts. Through open-source projects, it produces resources aimed at making the web more secure. More information is available at: https://owasp.org/
[14] The Industrial Internet Consortium (IIC) is a partnership of more than 200 corporations, as well as academic and research institutions - including Dell, Huawei, Microsoft, Kaspersky, and MITRE - dedicated to developing good practices for building trustworthy and reliable industrial IoT systems. More information is available at: https://www.iiconsortium.org/

Importantly, IIC also documents failure cases demonstrating how poor lifecycle management can produce severe security weaknesses (Buchheit et al., 2020).

- The Cloud Security Alliance[15] (CSA) provides further best practices for securing cloud-service architectures, including guidance emerging from leading cloud providers.

Based on the SDL experiences, Figure 4 depicts how security is connected to each stage of product development, demonstrating that secure-by-design is not a discrete step but a continuous process. While many development lifecycles formally end at release or at ongoing operations, several SDL models emphasise the importance of an *end-of-life* phase, in which vendors plan for de-supporting, retiring, or decommissioning products in a manner that minimises residual security risk. Some professionals suggest introducing the concept of *end-of-use* instead, highlighting that digital products often remain deployed – and therefore vulnerable – long after vendors stop issuing updates (Radunović and Grätz-Hoffmann, 2020). This underscores the need for suppliers to manage EOL transitions responsibly, communicate de-support timelines transparently, and provide guidance to users to reduce exposure from orphaned products.

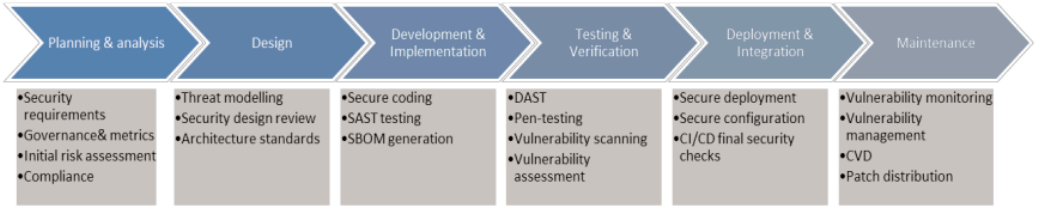


Figure 4: SDL mapped onto the standard digital product development process.

SDL models from the industry practice demonstrate a clear consensus: secure development must be embedded from the earliest design decisions onwards. Yet, despite the emphasis on preventing vulnerabilities from emerging in the first place, no development process can eliminate them entirely. This is why SDL frameworks increasingly converge on the need for robust vulnerability management processes – including CVD – as a necessary complement to secure design. The next section therefore turns to industry practices for vulnerability management and disclosure, examining how suppliers identify, assess, and mitigate vulnerabilities already present in deployed products.

3.4 Vulnerability Management and Coordinated Vulnerability Disclosure

Dominant industry frameworks – such as the Secure Development Lifecycle (SDL) – incorporate not only measures aimed at preventing or eliminating vulnerabilities during the design phase, but also robust processes for their management. Vulnerability management is a

[15] CSA brings together hundreds of technology companies – including Huawei, Google, Microsoft, IBM, and Oracle – to share experience and develop good practices for securing cloud services and related digital ecosystems. More information is available at: https://cloudsecurityalliance.org/research/artifacts/

core component of digital product security, given that software development processes unavoidably introduce defects (Spring, 2023). It encompasses the set of practices and security controls through which suppliers ensure that their products remain updated with relevant security fixes and mitigations over time (Radunović and Grätz-Hoffmann, 2020). In reality, it includes tasks like testing new products for vulnerabilities, scanning products that might be affected by discovered vulnerabilities, analysing and mitigating the effects of vulnerabilities in TPC integrated within own products, and producing, testing, distributing and deploying fixes and patches to minimise risks they pose (Householder et al. 2017).

Coordinated vulnerability disclosure (CVD) is of critical importance for vulnerability management. CVD refers to a structured process in which information about discovered vulnerabilities is shared and mitigated through coordinated action among cybersecurity researchers, manufacturers and vendors, possible intermediaries (like CERTs) and other interested stakeholders. In this process, parties work to agree on the timeline and conditions for publicly disclosing the vulnerability. The process provides the vendor with a defined window of time to develop and distribute a patch before the vulnerability becomes public. The international association of CERT/CSIRT teams, FIRST, has issued guidelines and good practices for such coordinated, multi-party processes (FIRST, 2020a), which become especially complex when vulnerability propagates through the supply chain.

The term responsible vulnerability disclosure is often used – particularly in non-technical or policy environments such as the UN cyber negotiations, where norm 13(j) of the 2015 UN GGE calls on states to encourage the 'responsible reporting of ICT vulnerabilities' (UN, 2015). Although frequently treated as a synonym for CVD, the term carries an ethical connotation and is less precise, as 'responsibility' may be interpreted differently by various stakeholders. From the vendor's perspective, responsible disclosure implies that researchers report a discovered vulnerability directly to them and provide sufficient lead time (commonly 15–45 days) to develop and deploy a patch before any public release of information, thereby reducing the likelihood of exploitation. Conversely, researchers understand responsible behaviour to include vendors' willingness to cooperate, communicate transparently, and act swiftly in producing a fix (Radunović and Grätz-Hoffmann 2020). At a minimum, the concept requires all parties involved to take all reasonable steps to minimise risks to users.

CVD, however, depends on the motivation of cybersecurity researchers. Böhme (2006) outlined several possible models through which researchers may monetise discovered vulnerabilities, including bug bounty programmes, vulnerability brokers, exploit derivatives, and cyber-insurance mechanisms. Miller (2007), however, highlights a number of challenges researchers face when attempting to sell vulnerabilities to legitimate buyers such as states or companies: discovered vulnerabilities have a limited shelf-life (i.e. until public disclosure, which may occur at any time); there are no established pricing models, as the value of a vulnerability depends on numerous factors and characteristics of its potential exploitability; it is nearly impossible to demonstrate a vulnerability to a buyer without revealing sensitive details that would allow the buyer to rediscover it independently; and it is difficult to enforce any

notion of ownership over a vulnerability, since it is essentially information that remains known to all parties once exchanged.

Not all researchers are motivated by money: many are ready to share their findings without financial rewards. Thus, effective vulnerability monitoring depends on visibility across multiple information sources. Empirical studies show that early indicators of newly discovered vulnerabilities frequently surface on social media platforms such as Twitter (now X) before they appear in vendor advisories, with dissemination often amplified through reposting and online discussion (Sauerwein et al., 2018; Syed et al., 2018). Building on this, Syed (2020) argues for systematically combining intelligence obtained from social media with structured vulnerability databases and proposes an ontology to manage and correlate these data. In parallel, underground and semi-public hacker forums constitute another important source of early signals; Benjamin et al. (2015) demonstrate how machine-learning techniques can be employed to collect, filter, and analyse such content to extract actionable vulnerability information.

There are also obstacles in the vulnerability disclosure process, ranging from the levels of trust of cybersecurity researchers in intermediaries or vendors to whom vulnerabilities are to be disclosed, to the surrounding regulatory and legal frameworks. Ethical researchers must be sure that disclosure will not trigger a legal process against them. They also expect timely, professional reaction from all the partners towards fixing the vulnerabilities (Walshe and Simpson, 2020).

In response, Herr (2017) advocates improving the legislative framework to incentivise researchers to disclose vulnerabilities, protecting them from legal processes, accelerating patch-development processes through greater vendor transparency about patch timelines, and increasing the number of user organisations that apply patches promptly. Involvement of accredited and trusted mediators can also help: national CERTs are playing an increasingly prominent role as intermediaries between researchers and suppliers, particularly in cases involving vulnerabilities with potential impact on critical sectors (Householder, 2020). Within the European context, for instance, recent regulatory developments further formalise this role by strengthening legal protections for security researchers and assigning CERTs explicit responsibilities in ensuring that vulnerabilities are remediated before public disclosure (Vander Geeten, 2024). This evolving legal landscape not only reduces the liability risks associated with vulnerability reporting but also creates stronger incentives for suppliers to cooperate with external stakeholders as part of a CVD process.

3.4.1 Industry practices for addressing vulnerabilities

Managing vulnerabilities in digital products requires suppliers to operate a complex set of interlinked technical, organisational, and collaborative processes. At a minimum, vendors must test new products for potential flaws; continuously monitor public reporting platforms and social media; receive, validate, and aggregate reports; scan products and track TPC

dependencies; assess severity; reflect on the required resources; develop, test, and deploy patches in a timely manner – tasks which demand extensive coordination across multiple internal units (security engineering, software development, quality assurance, product management, legal, and communications) as well as structured engagement with external actors such as OSS communities, security researchers, bug-bounty intermediaries, CERT/CSIRTs, downstream integrators, and end-users (Householder et al., 2020; Devaux et al., 2021; Radunović et al., 2021). Without this multilayered coordination, even technically robust fixes may fail to prevent exploitation in the wild.

Insights acquired through the multistakeholder dialogue of global industries, technical community, academia and civil society – the Geneva Dialogue on Responsible Behaviour in Cyberspace[16] – demonstrate that leading global manufacturers increasingly converge around three overarching categories of good practice for reducing vulnerability creation and exploitation:

- supply chain and TPC security
- secure development and deployment
- vulnerability management and disclosure (Radunović and Grätz-Hoffmann, 2020)

These pillars correspond to distinct stages of the product lifecycle, but in practice they operate as a continuous security loop, where lessons from incidents feed back into design, supplier governance, and development practices.

3.4.1.1 Supply Chain and Third-Party Security

Modern software development relies extensively, across the entire lifecycle, on third-party components (TPCs), including cloud services, open-source libraries and embedded firmware. This dependency chain exposes the industry to vulnerabilities they did not create, do not directly control, and are most often not aware about in the first place. Manufacturers often lack insight into the security practices of upstream vendors, and the absence of consistent certification frameworks makes it difficult to verify that a TPC meets an appropriate security standard.

This opacity complicates the detection of vulnerabilities in TPCs and creates systemic challenges: a single vulnerability in a widely reused library can propagate across hundreds of products and sectors, sometimes without the knowledge of the downstream integrator. These systemic, cascading risks are a major barrier to securing digital products and require visibility across the supply chain, and coordinated, multi-actor approaches (Kazakova et al., 2023).

TPC lifecycle which follows the SDL approach includes maintaining the inventory of all components including the OSS, accessing security risks, mitigating (or accepting) identified risks, and monitoring the changes in risk profiles of TPC (SAFECode, 2017). To put this into

[16] Available at: https://genevadialogue.ch

practice, the Software Bill of Materials (SBOM) has emerged as a practical mechanism to track components, versions, and dependencies, enabling suppliers to identify whether a newly disclosed vulnerability affects their products. However, tracking vulnerabilities throughout the supply chain requires resources and capabilities. As result, SBOM implementations vary considerably: some provide granular dependency graphs and license metadata, while others lack sufficient detail to support vulnerability detection (O'Donoghue et al., 2024). For SBOMs to become fully effective, they need to be integrated into automated pipelines that continuously update dependency information, aggregate information about related vulnerabilities and real-time threat intelligence, and trigger alerts when upstream changes occur (Shukla, 2025).

Big vendors have put this into practice in different ways. For instance, Microsoft integrates TPC visibility into its SDL, with SBOM to maintain a list of all third-party components, and the established plans to respond when upstream vulnerabilities emerge. Cisco maintains a Central Repository of Intellectual Property tracking all products using TPCs, enabling rapid identification of affected devices when a vulnerability in a dependency is discovered. Cisco also issues unique IDs for each produced device to prevent tampering and to ensure traceability through the supply chain. Others, like BI.ZONE, promotes using automated dependency-checking tools such as OWASP Dependency-Check[17] or Checkmarx Open Source Analysis[18] to detect insecure OSS components early in the development phase. (Radunović and Grätz-Hoffmann, 2020).

Beyond identifying components, suppliers must continuously monitor their TPCs for newly disclosed vulnerabilities – in proprietary and OSS components, but also in those that have reached end-of-life (EoL) – rapidly notify the TPC suppliers according to the best practice recommendations for cross-vendor cooperation, and ensure timely remediation. To achieve this, pre-emptive measures are also needed when making business agreements with suppliers. Leading companies have developed supplier-risk evaluation frameworks that include mandatory security requirements for component providers, contractual obligations for vulnerability notification, and periodic audits. Some firms also implement provenance controls, such as cryptographic signing of dependencies, secure software repositories, and sandboxing of untrusted components – measures which are particularly important for operational technology (OT) and embedded systems, whose long lifecycle makes them dependent on suppliers' long-term security commitments. (Radunović and Grätz-Hoffmann, 2020).

Finally, transparency is a fundamental enabler of supply chain security. Customers – especially organisational clients – often lack insight into the architecture, process flow, and third-party dependencies of products they deploy. Manufacturers must therefore adopt mechanisms that facilitate transparent communication about vulnerabilities, dependencies, and risks. SBOM helps customers understand the 'ingredients' of software and hardware they deploy, enabling informed risk analyses and procurement decisions. Transparency Centers[19] are another practice

[17] Available at: https://owasp.org/www-project-dependency-check/
[18] Available at: https://checkmarx.com/wp-content/uploads/2016/02/OSA-Datasheet.compressed.pdf
[19] Available at: https://www.kaspersky.com/transparency-center

embraced by the industry. For instance, Kaspersky Transparency Centers provide tiered access to insights about product engineering, security architecture, and TPCs, helping customers verify vendor claims and validate security measures in a controlled environment; Huawei Cyber Security Transparency Centre (HCSTC)[20] offers customers source code and binary testing capabilities (white-box and black-box), enabling independent validation during procurement and integration stages.

3.4.1.2 Secure Development and Deployment

A core pillar of security-by-design is embedding security controls, testing routines, and validation mechanisms throughout the development and deployment pipeline. As discussed earlier, security should not be treated post-festum but rather embedded into every phase of product development – from architecture design and coding to release, integration, and post-deployment operations. Modern development environments, especially those relying on automated continuous integration/continuous delivery (CI/CD) pipelines, now integrate mandatory security rules, static and dynamic analysis tools, dependency checks, and security-oriented unit tests to ensure that vulnerabilities are systematically identified and eliminated before products reach users (Allam, 2022).

Testing for software vulnerabilities and validation typically involves:

- Static Application Security Testing (SAST) remains a foundational technique in this regard. SAST tools analyse source code prior to compilation to identify unsafe functions, buffer manipulation issues, integer overflows, tainted inputs, and other defects that attackers frequently exploit.
- Dynamic Application Security Testing (DAST) complements SAST by examining the software in its compiled, running state. DAST tools simulate adversarial behaviour—probing interfaces, injecting malformed inputs, and triggering execution flows—to reveal memory corruption issues, privilege escalation opportunities, logic flaws, and other vulnerabilities only observable during runtime.
- Software Composition Analysis (SCA) serves for identifying and analysing the code of the secure third-party components, including the OSS, to detect security vulnerabilities, license issues, and compliance risks. SCA is of particular importance for generating SBOM, and responding to TPC risks. SCA tools can be integrated into the CI/CD pipeline, to enable continuous monitoring and testing in support of the security-by-design approach. (Marandi et al., 2023; Gaurav, 2024).

Companies increasingly integrate these mechanisms directly into developer workflows. For instance, Microsoft[21] deploys SAST throughout the pipeline to scan the code at rest, without running it, which allows developers to identify vulnerabilities early and act upon them before the build can proceed (Ayers, 2024). Cisco uses DAST as a black box testing for web

[20] Available at: https://www.huawei.com/-/media/CORPORATE/PDF/trust-center/huawei-cyber-security-transparency-centre-brochure-en
[21] Available at: https://www.microsoft.com/en-us/securityengineering/sdl/practices/security-testing

applications and APIs, building on community guidance like OWASP (Ramamoorthy, 2021). Various SCA tools emerged which are widely used due to their easy implementation and scalability; yet, number of them face challenges with accuracy of analysis and missing dependencies on vulnerable components (Jens et al., 2024). These mechanisms enable defenders to identify vulnerabilities introduced during product development, configuration or integration of TPC.

However, secure development does not end with testing individual components. Integrated product suites, systems-of-systems, and complex embedded environments require rigorous pre-deployment and post-deployment validation to identify vulnerabilities emerging from configuration choices or product interactions. Even well-tested modules can generate systemic weaknesses when assembled in particular network environments, making integration testing an indispensable step. Several industry frameworks (e.g. Microsoft SDL, NIST SP 800-218[22]) recommend documenting default configurations, validating them against common security baselines, and ensuring that insecure settings are eliminated before release. For critical infrastructure deployments, independent certification or repeat audits may be needed, though cost, difficulty maintaining audit currency for frequently updated software, and the absence of globally accepted certification standards for digital product security, remain the challenges (Radunović and Grätz-Hoffmann, 2020; Kazakova et al., 2023).

3.4.1.3 Vulnerability Processes and Support

Effective vulnerability management encompasses the full lifecycle of discovering, analysing, prioritising, and mitigating weaknesses in digital products. In industry practice, this requires a structured set of internal processes and close coordination with a range of external actors. Importantly, this understanding aligns with established definitions formulated by the Carnegie Mellon University CERT® (Householder et al. 2017), which underpin the ISO/IEC 29147:2014 standard on vulnerability disclosure (ISO/IEC, 2014). Together, these frameworks anchor contemporary corporate approaches to vulnerability reporting, management, and coordinated disclosure.

Suppliers must maintain reliable mechanisms for receiving vulnerability reports, validating and reproducing the issues, assessing severity, allocating resources, and deploying mitigations. Leading vendors have formalised these expectations through publicly available policies – such as the Cisco Security Vulnerability Policy[23] and Kaspersky Responsible Vulnerability Disclosure[24] guidance – which define intake channels (points of contact), communication timelines and procedures, and principles for coordinating disclosure with external parties (CVD policies). These policies aim to balance two imperatives: mitigating exploitation risk as quickly as possible and ensuring that disclosure processes protect users and do not inadvertently aid malicious actors.

[22] Available at: https://csrc.nist.gov/pubs/sp/800/218/final
[23] Available at: https://tools.cisco.com/security/center/resources/security_vulnerability_policy.html
[24] Available at: https://www.kaspersky.com/blog/vulnerability-disclosure-ethics/35581/

Bug-bounty programmes have become an integral element of many vendors' vulnerability management strategies. When properly scoped, authorised, and incentivised, they provide a structured mechanism for incorporating external cybersecurity research into organisational risk-management processes. Vendors may operate their own bug-bounty schemes – often rewarding researchers through non-financial incentives such as public recognition (e.g. 'hall of fame' listings) – but partnering with specialised third-party bug-bounty platforms that offer financial remuneration generally yields stronger results: empirical studies demonstrate that financial offers and prizes stimulate researchers and surge reporting (Zhao et al. 2015), broadening the pool of discoverable flaws and supporting earlier detection. Community-based disclosure platforms also generate valuable findings at comparatively low cost, functioning as an efficient complement to in-house engineering capacity (Ruohonen and Allodi, 2018; Walshe and Simpson, 2020). Yet these programmes are not without limitations: Sivagnanam et al. (2021), examining the large programmes like Chromium reward scheme, show that external researchers often uncover qualitatively different vulnerabilities from those exploited in the wild. Furthermore, Ruohonen and Allodi (2018) show that bug-bounty programmes do not necessarily reduce patch-delay challenges.

Externally, therefore, suppliers must cooperate closely with external partners, especially in case of CVD: security researchers, TPC maintainers, OSS communities, bug-bounty platforms, national or industry CERTs, CSIRTs and Information Sharing and Analysis Centers (ISAC), and customers and the public for issuing notifications and security advisories and distributing patches (Householder et al. 2020). Each actor contributes information that shapes how a vulnerability progresses from initial discovery to remediation, and each delay or failure in this chain increases the likelihood that attackers may weaponize and exploit the flaw.

Internally, vulnerability management is highly interdependent: security teams conduct initial scanning of code and TPC, conduct triage and technical validation; developers investigate root causes and prepare patches; legal teams manage liability assessments and regulatory obligations; communication teams prepare advisories and coordinate public notifications; C-level management monitors overall impact to business (Radunović et al., 2021). Vulnerability management, therefore, needs be a part of the organisational structure: manufacturers should establish the right mindset throughout an organisation because it is everyone's task (Charter of Trust, 2020).

Effective mitigation depends not only on technical capabilities but also on the organisational structures that coordinate and operationalise these responses. Many leading vendors institutionalise these responsibilities through dedicated product-security teams – typically PSIRTs – which serve as the central node for vulnerability triage, cross-departmental coordination, engagement with researchers and external partners, and communication with customers. The next section examines specifically the role and functioning of PSIRTs in more detail and discusses how they contribute to reducing the exploitability of vulnerabilities in real-world conditions.

3.4.2 Product Security Incident Response Team (PSIRT)

Suppliers' growing obligations – arising from customer expectations, regulatory requirements, and supply chains which are becoming very complex – increased demands for structured approaches to vulnerability response. As highlighted in the previous section, many regulatory frameworks and good-practice initiatives now explicitly recommend that vendors institutionalise dedicated product-security functions. In both academic and professional literature, the Product Security Incident Response Team (PSIRT) has emerged as a central model for organising these responsibilities (Rajnovic, 2010). Carnegie Mellon's CERT recommends PSIRTs as specialised teams focused on product-security incidents, capable of managing the CVD process, and distinct from enterprise-focused CERTs (Householder et al., 2020). ISO/IEC 30111 formalises this arrangement within its standard for vulnerability-handling processes and suggest PSIRTs are expected to coordinate internal mitigation activities while acting as the central interface for external stakeholders throughout the vulnerability disclosure process (ISO/IEC, 2019). Collectively, these frameworks depict the PSIRT as a core governance mechanism to ensure coherent internal processes and a trusted interface with external reporters.

3.4.2.1 PSIRT Tasks and Roles in CVD

A PSIRT is distinguished from CERTs and Security Operation Centres (SOCs) by its focus on suppliers' products and services rather than corporate IT infrastructure, although large organisations may operate both types of teams concurrently (Tanczer et al., 2018). At its core, a PSIRT is responsible for managing security issues across the entire product lifecycle. This, however, includes a variety of tasks:

- establishing transparent vulnerability-disclosure policies
- providing guidance for testing requirements and lessons-learned back to engineering
- developing incident-response procedures tailored specifically for product vulnerabilities
- maintaining communication channels with external actors such as researchers, industry peers, open-source communities and CERTs
- tracking vulnerabilities in TPC, using publicly available data and tools
- acting as the primary contact point for coordinated vulnerability disclosure (CVD)
- coordinating internal sectors and teams
- assessing internally discovered and reported vulnerabilities and conducting risk analysis
- supporting development and testing of fixes
- issuing advisories and customer notifications
- assisting with patch distribution
- promoting internal training and awareness (FIRST, 2019; Kasuya et al., 2023; Radunović and Grätz-Hoffmann, 2020; Rajnovic, 2010)

When a vulnerability is reported externally – whether by an independent researcher, a bug-bounty programme, or through an intermediary like a CERT – the PSIRT coordinates the entire response cycle. It receives the initial report, validates the technical findings, evaluates severity, identifies affected environments, determines the necessary mitigation. During the CVD process, the PSIRT negotiates disclosure timelines, ensures that patches or mitigations are prepared and tested, publicly discloses a vulnerability along with the advisory and a fix, and communicates clearly with researchers and affected stakeholders. Main PSIRT duties within the (ideal) CVD process are presented in Figure 5 (FIRST, 2019; Householder et al., 2020; Radunović and Grätz-Hoffmann, 2020; Siemens, n.d.). While the ideal CVD process follows a structured sequence, real-world responses may diverge due to uneven cooperation from upstream suppliers, inexperienced reporters, legal restrictions, or other external constraints. Nevertheless, the PSIRT remains responsible for steering the process towards transparent, risk-reducing outcomes (FIRST, 2019; Householder et al., 2020).

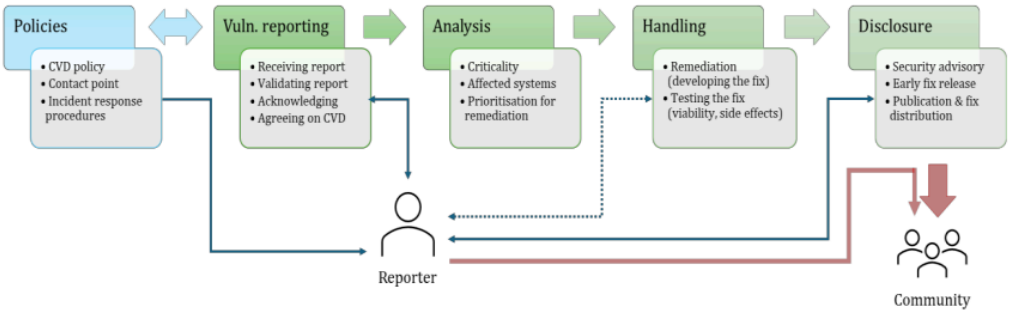


Figure 5. PSIRT role in a CVD process.

Several major suppliers illustrate how PSIRTs operate in practice mitigation (Radunović and Grätz-Hoffmann, 2020). Cisco's PSIRT[25], for example, integrates tightly with the company's secure development lifecycle (SDL). Its vulnerability-testing regimen includes protocol robustness testing, attack-simulation suites, and MITRE ATT&CK-aligned test cases, with triage performed jointly by product teams and PSIRT analysts to accelerate remediation. Siemens' PCERT[26] acts as a centralised contact point for vulnerability reports, oversees investigation – including for vulnerabilities originating in third-party components – and issues public advisories and mitigations. Adobe's PSIRT[27] similarly coordinates all aspects of vulnerability response across its extensive product suite. Across these examples, PSIRTs monitor disclosed vulnerabilities and exploits affecting proprietary and third-party components, exchange vulnerability information with trusted partners, and orchestrate internal mitigation and external communication to minimise the organisation's exposure window.

[25] Available at: https://www.cisco.com/c/dam/en_us/about/security/psirt/Cisco-PSIRT-Infographic.pdf
[26] Available at: https://new.siemens.com/global/en/products/services/cert.html#SecurityPublications
[27] Available at: https://www.adobe.com/content/dam/acom/en/security/pdfs/adb-incident-response-overview.pdf

3.4.2.2 Establishing and Running a PSIRT

Establishing and operating an effective PSIRT presents a range of organisational, legal, and technical challenges. One fundamental obstacle concerns adequate resourcing: PSIRTs require to work with multidisciplinary personnel – security engineers, incident responders, product architects, legal advisers, and communication specialists – yet many suppliers struggle to allocate stable budgets or hire staff with the required cross-domain expertise (FIRST, 2019; Tahara et al., 2025). PSIRTs increasingly require expertise in secure development, reverse engineering, threat modelling, supply-chain analysis, and incident communication – competences rarely found within a single individual and difficult to maintain without continuous training and organisational support (Cybellium, 2023).

PSIRTs must mediate between internal units that operate with different professional logics and priorities: for instance, Tahara et al. (2025) report friction between product safety teams – accustomed to conservative, certification-driven processes – and security teams that must act rapidly to address emerging threats, revealing persistent cultural and cognitive gaps that impede coordination. Geographical dispersion of development and operational teams further complicates coordination, especially in large multinational suppliers. Communication delays, inconsistent tooling, and uneven awareness of security expectations can hinder timely triage and remediation. Legal and regulatory pressures introduce further complexity, as handling sensitive vulnerability information entails non-negligible liability risks, particularly when products are deployed globally and obligations vary across jurisdictions (FIRST, 2019).

Vendor practice also shows that PSIRTs must coordinate with a diverse ecosystem of external stakeholders, each governed by different incentives and constraints. Supply-chain partners may have heterogeneous security maturity levels; open-source communities are decentralised and largely voluntary; researchers and bug-bounty participants vary widely in expertise and expectations; CERTs operate under national mandates; and bug-bounty platforms have commercial interests that shape communication flows. Managing these relationships demands not only technical competence but also diplomatic skill, transparency, and clearly formalised processes (Householder et al., 2020; Cybellum, 2023). As the FIRST PSIRT Maturity Model notes, organisations often advance from ad-hoc practices to structured, policy-driven coordination only gradually, and many suppliers remain at early to intermediate stages – lacking comprehensive workflows, cross-team training, or mechanisms for systematically documenting and learning from past incidents (FIRST, 2019).

3.4.2.3 Challenges for Establishing PSIRT

Taken together, the organisational, technical, and collaborative demands associated with establishing a PSIRT demonstrate that this function represents far more than an internal procedural adjustment. Instead, it constitutes a substantive organisational transformation that requires sustained investment, senior-level commitment, and cross-functional coordination. Consequently, the capacity to operate a mature PSIRT is unevenly distributed: it is

disproportionately feasible for large firms headquartered in technologically advanced economies, or for companies that explicitly prioritise cybersecurity maturity as a strategic capability.

To examine how PSIRTs are distributed across the global digital-product ecosystem, a dataset was compiled of suppliers maintaining publicly identifiable PSIRTs registered with key international coordination mechanisms – namely the FIRST[28], the Trusted Introducer framework for TF-CSIRT teams[29], and the CVE Numbering Authority (CNA) programme[30]. These three mechanisms were selected because participation in them requires demonstrable organisational maturity, established disclosure processes, and the ability to engage in transnational information exchange on vulnerabilities. Although this dataset does not capture all existing PSIRTs – particularly those operating internally without public affiliation – it provides a reliable sample of actors capable of participating in the global vulnerability-coordination architecture. After removing duplicates and validating organisational identities, the resulting dataset contains approximately 80 companies Their distribution across geography, company size, and industry sector reveals several instructive patterns.

First, more mature PSIRTs are overwhelmingly concentrated in a small number of jurisdictions. More than half of the identified PSIRTs belong to companies headquartered in the United States, followed by EU member states (notably Germany, Sweden, Finland, France, and the Netherlands), and then Japan, Taiwan, China, and Russia (Figure 6). This suggests a strong correlation between states with strong digital-product industries and vendors with institutional capacity for PSIRTs.

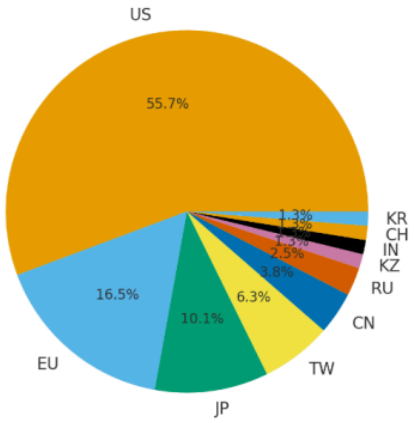


Figure 6: Geographical distribution of identified PSIRTs. This figure shows the proportion of PSIRTs associated with companies headquartered in different geographies, with dominance by the United States, the European Union, and other regions (Japan, Taiwan, China, Russia). Data source: author's dataset (2025).

[28] Available at: https://www.first.org/members/map
[29] Available at: https://www.trusted-introducer.org/tf-csirt-community/
[30] Available at: https://www.cve.org/partnerinformation/listofpartners

Second, PSIRTs are predominantly a feature of large enterprises. Nearly 20% of identified PSIRTs belong to 'mega-firms' with workforces exceeding 100,000 employees, and a further 20% belong to very large companies (50,000–100,000 employees). Interestingly, however, almost half are associated with large companies employing 1,000–50,000 staff. Only slightly more than 10% are found in small and medium-sized enterprises (SMEs) (Figure 7). This distribution highlights the resource-intensive nature of PSIRT operations, suggesting that smaller firms face notable barriers in adopting comparable practices.

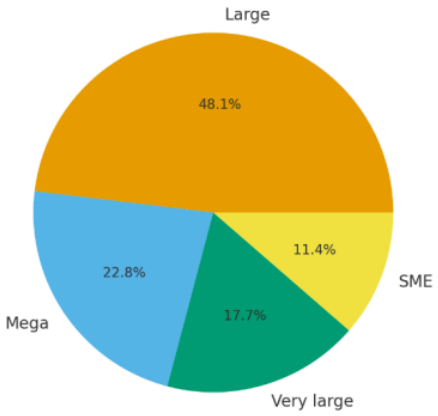


Figure 7: Distribution of PSIRTs by company size. Companies are categorised by workforce (SME <1,000; Large 1,000–50,000; Very large 50,000–100,000; Mega >100,000 employees). Data source: author's dataset (2025).

Third, sectoral patterns indicate that more than half of all PSIRTs belong to software vendors, reflecting the centrality of software supply chains in modern security ecosystems. Encouragingly, over one quarter are maintained by suppliers serving highly sensitive or regulated industries – such as energy, telecommunications, and healthcare – where vulnerability exploitation carries disproportionately high societal risks. By contrast, cybersecurity vendors represent only around 5% of the identified PSIRTs (Figure 8), an unexpectedly low figure given their central role in vulnerability research and incident response.

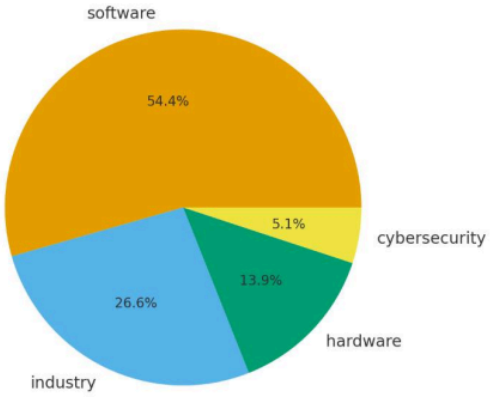


Figure 8. Distribution of PSIRTs by industry sector. Sectors include software, hardware/IT, cloud services, industry/manufacturing, telecommunications, healthcare, cybersecurity, and others. Data source: author's dataset (2025).

Taken together, these findings show that PSIRT has become a recognised good practice among the world's largest technology producers, yet its adoption is uneven and strongly correlated with organisational size, financial capacity, and geopolitical location. Awareness about this good practice and its merits also plays a role. This reinforces the conclusion that PSIRT implementation, while regarded as a good practice in international policy and industry guidance, demands significant financial, technical, and organisational capacity.

Despite these challenges, the PSIRT model offers a coherent and comprehensive structure for managing the full spectrum of vulnerabilities associated with digital products. It integrates and aligns multiple domains – secure development, supply chain and TPC risk management, and vulnerability management practices – under a single organisational framework.

This makes the PSIRT particularly relevant for the purposes of this study: as a consolidated practice, it allows for a more straightforward assessment of the effectiveness of suppliers' internal processes in reducing the likelihood of vulnerability exploitation in the wild, and thus, indirectly, on constraining the operational freedom of OCCs. The following section therefore examines specific metrics and indicators that can be used to evaluate the efficiency and impact of such practices.

4 Assessing the Efficiency of Good Practices

The preceding sections establish SDL, structured vulnerability management, and dedicated PSIRTs as widely recognised good practices for improving the security of digital products. Conceptually and operationally, it is reasonable to expect that those practices contribute to greater security of digital products. For instance, it is reasonable to expect that a PSIRT – by institutionalising internal coordination and enabling sustained engagement with external security researchers, vendors, and coordination bodies – can reduce exposure to actively exploited vulnerabilities. In policy, standards, and industry discourse alike, PSIRTs are therefore commonly presented as a key mechanism for mitigating cyber risk.

Yet recognising a practice as 'good' does not, in itself, demonstrate its effectiveness. To substantiate such claims, it is necessary to assess and quantify whether and how these practices measurably influence vulnerability exploitation in practice. This requires moving beyond inputs and processes – such as the existence of policies, teams, or procedures – and examining outcomes, particularly the likelihood that reported vulnerabilities are exploited in the wild. Such a focus is crucial, as not all vulnerabilities are of interest to adversaries, and mitigation alone does not equate to reduced strategic risk, especially in the context of OCO.

Accordingly, this section reviews existing industry and academic methodologies used to assess the effectiveness of security-by-design practices, with a particular focus on PSIRT-related metrics. It then identifies key gaps in current approaches and distils a set of elements relevant to evaluating whether and to what extent good practices contribute to reducing real-world exploitation. These considerations form the analytical foundation for the methodology developed and applied in the following chapters.

4.1 Industry Methodologies

It is essential to verify the effectiveness of secure-development frameworks and good-practice methodologies as implemented by companies. A few vendors publish periodic reports or maintain internal metrics on their adoption of secure-development lifecycles, but these seldom assess whether the adoption actually leads to fewer or less severe real-world vulnerabilities.

Microsoft, in its early report on progress in the implementation and further development of the SDL, presents advances in both the evolution and components of its SDL framework, as well as the level of adoption of specific security elements by third parties developing applications for the Windows operating system (Ladd et al., 2011). Microsoft's model is also convenient for other companies, as it offers assessments their product-security maturity through so-called readiness 'scorecards', based on metrics derived from steps within the SDL framework. However, these reports do not evaluate the effectiveness of the SDL itself in relation to vulnerabilities – namely, whether the number and criticality of vulnerabilities in Microsoft systems and in third-party applications have declined because of SDL adoption.

Waidner, Backes and Müller-Quade (2014) cite a Microsoft report from 2013 claiming that the number of disclosed vulnerabilities within the first year after release dropped by 45% when comparing Windows XP to Vista, and by as much as 91% when comparing SQL Server 2000 to SQL Server 2005, attributing this reduction to the introduction of the SDL approach; the cited Microsoft report, however, is no longer publicly available. More recently, Microsoft reported it published over 1000 vulnerabilities – around 50 of which requiring no action – throughout 2025, out of which 48% were discovered in-house, while their use of AI in vulnerability triage contributed to a 72% success rate addressing them within the reduced time to mitigate (Microsoft, 2025). While this report tracks internal metrics (time-to-mitigate, number of triaged and published vulnerabilities), it doesn't show external outcome metrics (e.g. incident rate, patch-gap exploitation, comparative risk before vs after SDL).

A more systematic option comes with maturity-models that allow organisations to self-assess and benchmark their security posture. The OWASP Software Assurance Maturity Model (SAMM) v2[31] offers a flexible, risk-driven, lifecycle-agnostic framework that spans governance, design, implementation, verification, and operations. Its structure of business functions, security practices and maturity levels enables organisations – regardless of size or sector – to evaluate their current practices and set incremental targets for improvement. In a recent industrial case study leveraging SAMM, Fucci et al. (2024) found that different professional roles within the same organisation perceive maturity differently, but also that a lightweight SAMM-based assessment is viable and effective for quickly evaluating and improving security posture. There is, however, no metrics on how the OWASP maturity affects the exploitation of vulnerabilities.

The Building Security in Maturity Model (BSIMM) represents the most sustained industry-driven attempt to empirically observe how organisations implement secure development and product-security practices over time. Since 2008, BSIMM studies have analysed the adoption of approximately one hundred security practices across a growing cohort of companies, spanning software, cloud services, IoT, finance, healthcare, and industrial sectors. The BSIMM10 edition of the study from 2019 examined 122 organisations and structured practices across four domains – governance, intelligence, secure software development lifecycle, and deployment – focusing on the presence and maturity of practices rather than their direct security outcomes (Migues et al., 2020). The most recent iteration, BSIMM15 from 2025, expands the dataset further and confirms long-term trends: most organisations prioritise vulnerability response and remediation, followed by pre-release testing, while proactive architectural and design-level interventions remain less consistently adopted (Black Duck, 2025).

Importantly, BSIMM data emphasise that the model assesses process maturity and organisational behaviour, not reductions in exploited vulnerabilities or incidents. A study of BSIMM data over 12 years, conducted by Weir et al. (2022), reinforces this limitation, noting that BSIMM enables benchmarking and peer comparison but does not establish causal links between specific practices and measurable security outcomes such as exploit prevalence or

[31] Available at: https://owaspsamm.org/assessment/

incident reduction. As such, BSIMM remains a valuable descriptive and comparative instrument for understanding how security practices diffuse across industry, while highlighting a persistent methodological gap in evaluating their real-world effectiveness against vulnerability exploitation.

As potential criteria for assessment of the efficiency of good practices, industry discussions also suggest process maturity (e.g. reaction speed, clarity and transparency of procedures, extent of coordination with industry and CERTs), as well as quantitative parameters such as time-to-fix, change in number or severity of vulnerabilities discovered over time (internally or externally), and patch-deployment rates. Using a common denominator like the CVSS may help standardise such assessments across companies and time (Radunović and Grätz-Hoffmann, 2020).

However, optimising internal processes to score highly on maturity or speed does not – individually – guarantee better security in practice: faster patching may come at the cost of incomplete fixes; widespread adoption of security measures does not, in itself, prevent novel vulnerabilities. Ultimately, to claim that a given practice or set of practices genuinely reduces the likelihood of exploitation – and thereby contributes to deterring offensive cyber operations – one needs empirical studies that correlate (and ideally demonstrate causation between) process maturity, vulnerability incidence/severity, exploit prevalence, and security outcomes. As of now, this remains largely an underdeveloped area of research.

4.2 Academic Methodologies

Academic research has approached the question of effectiveness of vulnerability-related practices from several angles, most often focusing on disclosure dynamics, patching behaviour, or process efficiency, rather than on direct reductions in exploitation in the wild.

Effectiveness of vulnerability disclosure and patching practices is sometimes approached through empirical observation. Arora et al. (2004) conducted an early empirical study of 308 vulnerabilities, correlating vendor responses with exploitation data collected through honeypots designed to attract attacks. Their work analysed how quickly and effectively companies responded once vulnerabilities became known. Given the significant evolution of vendor practices, coordinated disclosure programmes, and regulatory pressures since then, these findings are best interpreted as historically informative rather than directly generalisable to current practices.

More recent work has examined efficiency through the lens of secure development and operational workflows. Ramaswamy (2023) proposes a development, security, and operations (DevSecOps) implementation that embeds automated security controls into a CI/CD pipeline and evaluates it through a simulated agile development scenario. The evaluation relies on operational metrics such as vulnerability detection rate, scanning latency, false-positive ratios, and mean time to remediation. While these metrics are valuable indicators of process efficiency

and tool performance, the study does not establish a link between such improvements and a reduced likelihood of real-world vulnerability exploitation.

A more targeted performance metric for vulnerability response was proposed by Frei et al. (2008). Comparing Microsoft's and Apple's handling of zero-day vulnerabilities, they introduced the zero-day patch rate as a measure of effectiveness – defined as the number of vulnerabilities patched by a vendor on the same day the vulnerability is publicly disclosed. This approach provides a concrete and comparable indicator of responsiveness, though it remains limited to observable patching behaviour rather than downstream exploitation outcomes.

Scholarly work that explicitly examines PSIRTs is still limited. The existing literature predominantly consists of organisational case studies or the development of technical tools aimed at improving PSIRT operations.

For example, Morrison et al. (2017) examined security practices within IBM's industrial software development teams and, drawing on qualitative analysis and personnel surveys, ranked vulnerability tracking – organised through the IBM PSIRT – among the three most consequential security measures. Subsequent research has largely explored how automation can enhance PSIRT operations. For example, Shourya et al. (2023) present an automated vulnerability management system designed to support PSIRT analysts and report a 56% improvement in task efficiency. Kasuya et al. (2023) demonstrate that an automotive-focused tool based on NVD data can eliminate up to 96% of irrelevant information, substantially reducing the analytical burden of PSIRTs. Although these contributions demonstrate gains in internal efficiency and workflow optimisation of PSIRTs, they stop short of assessing if such improvements have a measurable effect on vulnerability exploitation in real-world settings.

Because PSIRTs share several operational characteristics with incident response teams, the scholarly work on SOCs offers helpful methodological reference points. Existing studies consistently emphasise metrics related to vulnerabilities, such as the extent of exposure, time to remediation, severity assessments, and the occurrence of incidents tied to already-known flaws (Vielberth et al., 2020). A systematic review by Forsberg and Frantti (2023) similarly finds that vulnerability-focused indicators rank among the most commonly used measures of SOC performance, reflecting both their prominence in widely adopted standards like NIST and ISO/IEC and the fact that they can be quantified using available data. Survey-based research shows that organisations frequently rely on numerical indicators such as the number of incidents processed and average response times when evaluating SOC effectiveness (Crowley and Pescatore, 2019). Although these approaches can inform the development of metrics for PSIRTs, important differences remain: SOCs are primarily concerned with operational incident handling, whereas PSIRTs also encompass coordinated vulnerability disclosure and remediation activities that are tightly coupled to specific products.

Overall, while existing academic studies offer valuable insights into responsiveness, process maturity, and operational efficiency, the existing literature does not yet provide a comprehensive empirical framework for evaluating how PSIRT – and vulnerability

management practices more generally – influence the likelihood of the active vulnerability exploitation. That gap underscores the need for methodologies that connect organisational practices more directly to security outcomes and strategic risk.

4.3 Key Elements for Assessing PSIRT Efficiency

Building on the review of industry methodologies and academic studies, this section consolidates the key metrics and indicators that are commonly used – or proposed – for assessing the performance and maturity of product vulnerability handling practices, particularly PSIRTs. These elements arise from three main bodies of work: (i) industry SDL and maturity models, (ii) empirical and design-oriented academic studies on vulnerability management, and (iii) adjacent research on incident response and SOC performance. Taken together, they provide a broad but fragmented set of assessment criteria.

4.3.1 Overview of Commonly Used Metrics

Table 2 synthesises the principal assessment elements identified across the reviewed literature and practice, grouping them thematically and indicating their primary origin.

Category	Metric / Indicator	Rationale
Process maturity	Defined CVD policy and workflows	Indicates organisational readiness and institutionalisation
	PSIRT staffing and roles	Indicates organisational capacity
Process efficiency	Mean time to triage	Captures responsiveness to incoming reports
	Mean time to remediate or patch	Widely used operational efficiency metric
	Patch issuance at disclosure (e.g. 0-day patch rate)	Directly linked to exposure window
Output volume	Number of vulnerabilities disclosed	Signals efficiency, visibility and reporting culture
	Number of advisories issued	Indicates disclosure activity
Severity-related	Distribution of CVSS scores	Enables comparison across products
Trend-based	Change in critical vulnerabilities over time	Captures potential structural improvements
External engagement	Bug-bounty participation	Expands detection surface
Coordination	CERT / CSIRT collaboration	Facilitates information sharing
Automation support	Use of automated triage tools	Improves analyst efficiency

Table 2: Synthesis of the principal assessment elements identified across the literature

4.3.2 Metrics with Limited Explanatory Power for Exploitation Risk

A substantial subset of these indicators measures process maturity or operational efficiency, rather than security impact. Metrics such as the existence of formal policies, internal staffing levels, or the speed of internal workflows are valuable for benchmarking organisational development, but they remain indirect indicators of security outcomes. As highlighted in both industry and academic critiques, optimising these metrics may sometimes generate counterproductive effects – for example, prioritising rapid patch release at the expense of testing quality or coordination.

Similarly, direct counts – such as the number of vulnerabilities disclosed or advisories published – are highly sensitive to detection capability, disclosure culture, and regulatory pressure. Counting can also prioritise volume over severity or particular risk of exploitation. This metrics cannot be interpreted without contextualisation and do not distinguish between benign disclosure dynamics and reduced attacker opportunity.

4.3.3 Metrics with Stronger Explanatory Power for Exploitation Risk

A smaller set of indicators shows clearer conceptual links to exposure and exploitation risk. These include:

- patch availability, which directly affects the window of opportunity for attackers
- vulnerability severity, when combined with exploitability signals
- coordination indicators, particularly where CERTs or multi-vendor disclosure processes reduce systemic exposure across supply chains
- time-bound remediation metrics (e.g. time to patch), assessed over time to capture trends and consistency rather than as single point measurements

The existence of a PSIRT does not, in itself, determine whether a patch or mitigation will ultimately be developed: remediation may be produced by internal development teams or, in the case of OSS components, by the wider community. Likewise, the severity of a vulnerability is independent of whether a supplier maintains a PSIRT. By contrast, indicators related to coordination, communication, and remediation processes – such as the timeliness, consistency, and transparency of disclosure and patch delivery – are directly shaped by the presence and functioning of a PSIRT.

Notably, even where such metrics are well defined and systematically collected, they are seldom examined in relation to observable patterns of vulnerability exploitation in the wild. As a result, they provide only limited insight into how organisational practices interact with external risk factors – such as the strategic value of a product, its widespread deployment, supply-chain reuse (including open-source components), or its attractiveness to state and non-state actors – which are central to understanding vulnerability exploitation in the context of OCO.

4.3.4 External Factors Relevant for Exploitation Risk

Assessing the level to which a PSIRT decreases the risk of real-world vulnerability exploitation requires accounting for external factors that lie beyond the control of the supplier. Of particular relevance to this research are cases of exploitation associated with state interests – situations in which vulnerabilities become instruments in OCO targeting suppliers or systems whose compromise may generate political, economic, or strategic advantage. In such contexts, malicious actors selectively target digital products whose disruption can produce significant downstream consequences across dependent systems and supply chains.

Therefore, evaluating PSIRT effectiveness cannot rely solely on internal process indicators but must consider whether vulnerabilities addressed by PSIRTs are those most likely to be exploited in practice. This entails focusing on vulnerabilities with a high interest for active exploitation, especially in contexts relevant to OCO and APT operations, where the targets are either intrinsically high-value or integrated into critical infrastructures that amplify the potential impact of compromise.

4.3.5 Foundations for a New Methodology

To conceptualise the risk of active exploitation, we return to the general risk formulation (1), and adapt it according to Lee (2014) to capture software-specific considerations such as architectural complexity and availability, integrity, and confidentiality requirements:

(2) $R = L \times S \times I$

where

- R represents the risk of active exploitation
- L represents the likelihood of vulnerability exploitation
- S represents the vulnerability severity
- I represent the impact of exploitation

Severity (S) is typically measured by the CVSS scores. While certain CVSS elements – particularly the confidentiality and integrity dimensions of the *Impact* sub-score – have some relevance to exploitation likelihood, this relevance is limited (Allodi and Massacci, 2014). Importantly, severity is independent of the existence of a PSIRT, as it reflects characteristics of an already-existing vulnerability to which a PSIRT responds, rather than an outcome of PSIRT activity.

Impact (I) captures the potential consequences of exploitation and depends on both technical characteristics of the vulnerability and external contextual factors. The technical dimension is reflected in CVSS impact metrics, already accounted for. Beyond this, impact is strongly shaped by the nature of the affected asset: the function of the vulnerable product, its role within wider ICT systems, and the societal or economic importance of the supplier. Vulnerabilities

affecting products deployed in critical environments – such as energy systems, financial services, or healthcare – can yield disproportionate societal harm. As with severity, the impact of exploitation is independent of whether a supplier maintains a PSIRT.

By contrast, the likelihood of exploitation (L) is directly linked to how vulnerabilities are managed – that is, how rapidly and effectively they are identified, assessed, mitigated, and disclosed. These processes lie at the core of PSIRT responsibilities. The availability of a patch or mitigation therefore emerges as an important indicator.

The likelihood of exploitation, however, is influenced by other discussed factors as well – notably the existence of proof-of-concept (POC) exploits, the availability of weaponised exploits, and their circulation on dark markets or offensive tooling platforms. Exploit and POC availability are external conditions shaped largely by the behaviour of researchers or attackers, thus not a factor of PSIRT presence.

Another factor is a type of a product or industry which is targeted: certain type of targets, like the widely used products or critical infrastructure suppliers, particularly attract actors with state-backed or geopolitical motives. Vulnerabilities in open-source components also merit attention: although OSS provides transparency, many projects lack resources for rigorous security review, and widely used components can become vectors for supply-chain compromise. In addition, the still-limited uptake of SBOMs across sectors heightens the difficulty of tracing vulnerable dependencies, leaving latent security gaps that sophisticated adversaries can exploit. Thus, OSS may also contribute to the likelihood. Neither a type of industry nor a type of product, however, depend on the presence of PSIRT.

Accordingly, assessing risk as a function of the presence of PSIRT as a good practice can be reduced to modelling the likelihood of vulnerability exploitation as a function of whether such a team is in place, taking into account the relevant indicators:

- availability of a patch or a mitigation
- public availability of exploits (e.g. dark markets or POC releases)
- supplier's sector (e.g. critical infrastructure, general IT, cloud platforms)
- openness of the product (OSS versus proprietary)
- product popularity (e.g. size of user base or deployment footprint)

Following the work by Radunović et al. (2025), the next section proposes a methodology for modelling this dependence, integrating the internal and external indicators. This in turn allows for a rigorous examination of how the adoption of PSIRT practices – particularly among leading global suppliers – may influence the broader landscape of OCO.

5 Proposed Methodology for Assessing the Efficiency of PSIRT

To assess whether PSIRT can indeed be considered an effective good practice, it is necessary to develop a methodology capable of evaluating its impact on the real-world exploitation of vulnerabilities and, by extension, on the feasibility and effectiveness of offensive cyber operations. Rather than relying on theoretical assumptions or process maturity alone, such an assessment must be grounded in empirical evidence on vulnerabilities that are disclosed and actively exploited in practice, and linked to the organisational practices of their suppliers, including the presence or absence of a PSIRT.

To this end, this study adopts a methodological approach well established in medical and epidemiological research and increasingly applied in cybersecurity studies: the matched case-control observational design. In this framework, the existence of a PSIRT is treated as an exposure factor, while other influential characteristics are controlled for. The analysis draws on trusted public datasets on disclosed and exploited vulnerabilities and on suppliers' vulnerability-handling practices. To reduce noise and mitigate potential biases, bootstrapping techniques are applied. This chapter details the matched case-control design, defines the variables relevant to this study and the resulting model structure, describes the datasets used, discusses methodological limitations, and explains the data preprocessing steps.

5.1 Matched Case-control Design

A matched case-control study is a form of analytical observational research commonly used in clinical and epidemiological settings. In this retrospective design, outcomes (such as a disease or the failure) are first identified and then compared across observational units that differ in their exposure to specific factors (such as risks or practices), while matching on other characteristics to control for influences not under direct examination that could otherwise affect the results.

In medical research, case-control designs are applied both in *prospective* and *retrospective* settings. Prospective studies typically assess outcomes following deliberate exposure of observational units to an intervention or risk factor, such as in clinical trials evaluating new pharmaceuticals. In many situations, however, deliberate exposure is impractical or ethically impermissible. As a result, *retrospective* case-control studies are used instead, examining existing outcomes to infer how particular risk factors may have contributed to them, rather than to test the effect of a planned intervention.

Although retrospective approaches are well suited for identifying correlation between risk factors and observed outcomes, they offer limited capacity to establish causality, owing to the absence of experimental control. Classic examples include studies on the relationship between smoking and lung cancer (Doll and Hill, 1950), and analyses of seatbelt effectiveness in reducing fatalities in traffic accidents (Evans, 1986). Within cybersecurity research, a comparable methodology has been adopted by Allodi and Massacci (2014), who used a case-

control framework to assess the contribution of CVSS scores to vulnerability exploitation risk, relative to the presence of proof-of-concept exploits and exploit availability on underground markets.

In a retrospective case-control study, observational units are selected based on their outcome status (e.g. a disease or an adverse event – or exploitation of a vulnerability in the wild, in our case): *cases* are samples with an outcome, while *controls* are those without the outcome. At the same time, exposure of observational units to certain factors (e.g. risk factors, treatments – or exposure to a good practice, in our case) is assessed retrospectively; information about the exposure is collected through existing records or interviews.

Therefore, observational units are organised according to:

1) exposure to a risk factor: *exposed* versus *unexposed*
2) outcome status: cases (with an outcome) versus controls (without an outcome)

We then count number of elements with each label:
a – the number of all exposed cases
b – the number of all exposed controls
c – the number of all unexposed cases
d – the number of all unexposed controls
This is presented in Table 3:

	Cases (with outcome/ adverse event)	Controls (without outcome/ adverse event)	
Exposed (risk factor, treatment)	*a*	*b*	Total exposed: *a*+*b*
Unexposed (risk factor, treatment)	*c*	*d*	Total unexposed: *c*+*d*
	Total cases: *a*+*c*	Total controls: *b*+*d*	

Table 3: Labelling of the observational units according to exposure and outcome status

In a matched case-control design, controls are selected to match cases on certain characteristics that may also influence the outcome (the confounding variables). This allows to eliminate the impact of confounding factors on the outcomes. Each case has a single matching control. Since, for each case, there is a single chosen control, the total number of cases needs to be equal to the total number of controls: $a+c=b+d$.

Yet, when the pool of available controls is larger than that of cases, we have more options to select a matching control for each case. In such case, the selection of a matching control for each case is done randomly, allowing us to do re-sampling multiple times. This technique is

known as *bootstrapping*: it flattens noise in the data by repeatedly resampling it with replacement, creating multiple datasets. Randomisation ensures that, with a sufficiently large sample, potential confounding variables – whether observed or unobserved – are equally distributed across groups. This balance prevents confounding variables from correlating with the independent variable, minimising their impact on the study. As a result, randomisation is widely regarded as the most effective way to address confounding.

Depending on the pool of available controls to select from, re-sampling can be done many times, and the median of the results can be considered for the analysis. Such a randomised matched case-control design reduces eventual noise in the data and isolates the effect of other external factors irrelevant to the study – including those that can't be directly observed.

5.2 Variables

To evaluate the efficiency of PSIRT in reducing the likelihood of exploitation, we should observe performance of suppliers of digital products: those that do have a PSIRT, and those that don't. Their performance, however, needs to be assessed on the level of vulnerabilities discovered in their products, since eventual exploitation happens to each individual vulnerability.

In this study, we observe vulnerabilities known to have been exploited in the wild, and their matching pairs for which there is no record of exploitation in the wild. We analyse the effects of PSIRT of the corresponding supplier on the likelihood of exploitation of a vulnerability in its product. Since we hypothesise that the good practice impacts the likelihood of exploitation in the wild, we consider a PSIRT put in place by a supplier as a *risk factor*.

5.2.1 Independent and Dependent Variables

Correspondingly, transforming the design from Section 5.1 for our case:

- *risk factor,* as an independent variable whose value we set, is a presence of a good practice – PSIRT, within a supplier of a vulnerable product
- *outcome,* as a variable dependent on risk factors, is an exploitation of a vulnerability in the wild
- *observational units* are known vulnerabilities in digital products, which are directly linked to a supplier of a vulnerable product; in particular:
    - *cases* are vulnerabilities which are known to have been exploited in the wild
    - *controls* are other vulnerabilities chosen at random from the population of vulnerabilities of suppliers that have no record of vulnerabilities exploited in the wild, and whose characteristics match those of the cases (based on the confounding factors and the similarity function discussed below)

We label each vulnerability *v* of the sample according to:

1) Dependent variable:
   a) it *is exploited* in the wild – thus belonging to a set of cases C, or
   b) it *is not exploited* in the wild – thus belonging to a set of matched controls M
2) Independent variable:
   a) its supplier *is exposed to PSIRT as a good practice* – thus belonging to a corresponding set G, or
   b) its supplier *is not exposed to PSIRT* – thus not belonging to a set G

We now adapt Table 3 for our study (Table 4):

	Cases: $v \in \mathbf{C}$ (vulnerability v exploited in the wild)	Controls: $v \in \mathbf{M}$ (vulnerability v not exploited in the wild)	
Exposed: $v \in G$ (supplier has a good practice, i.e. PSIRT)	$a = \lvert \{v \in C \ \& \ v \in G\} \rvert$	$b = \lvert \{v \in M \ \& \ v \in G\} \rvert$	Total exposed $(a+b)$
Unexposed: $v \notin G$ (supplier doesn't have a good practice, i.e. PSIRT)	$c = \lvert \{v \in C \ \& \ v \notin G\} \rvert$	$d = \lvert \{v \in M \ \& \ v \notin G\} \rvert$	Total unexposed $(c+d)$
	Total cases $(a+c)$	Total controls $(b+d)$	

Table 4: Labelling of the observational units according to exposure to PSIRT and the exploitation status

where $a+c=b+d$, since $|C|=|M|$ by design.

5.2.2 Confounding Variables

While focus of this methodology is how presence of a PSIRT within a supplier impacts likelihood of exploitation of vulnerabilities in the wild, there might be other influential factors. In our study, confounding variables are such factors that, combined with a presence of PSIRT, may correlate with the observations about the outcomes or be alternative explanations. Therefore, we must elaborate on possible confounding variables, and take the appropriate ones into account – but only if the related data is available.

Reflecting on earlier discussion about indicators and external factors that impact risk from exploitation, we conclude that the availability of POC for vulnerabilities can be considered as a confounding variable, since this information is publicly available. Information about the availability of exploits in the dark markets, however, is typically available only to malicious

actors that are members of such markets; due to lack of authoritative publicly available regularly updated database, we cannot account for it in our study.

Presence of a fix (a mitigation or patch) should also be considered a confounding variable. Information about mitigation options and patches are publicly available, and we can thus consider presence of a fix as a confounding variable.

The supplier's industry type is another important factor to consider in this study. Since this information is publicly accessible and can be systematically aggregated, it is possible to aggregate it and include it as a confounding variable. For our analysis, suppliers are categorised into the following groups: software, hardware, cloud services, industry and manufacturing (which often have higher criticality and may attract specific malicious actors), and cybersecurity products (to observe the performance of this sector). Collecting more granular data about each supplier – such as organisational maturity and regulatory obligations – would be challenging at this stage, and can be done as part of the future work. Information about open-source or proprietary nature of the products of a supplier is public and possible to collect, allowing us to consider whether a supplier delivers OSS as a confounding variable.

Malicious actors are often drawn to vulnerabilities in popular digital products with widespread use. While metrics like application download counts could serve as indicators, they are unreliable and do not account for digital products embedded in other systems. Since such information is neither publicly available nor systematically aggregated, we cannot include product popularity as a confounding variable.

Similarly, products used in critical environments (e.g. finance, energy, healthcare) may attract APT groups. However, apart from specialised products from critical industries already considered under the supplier's industry type, it is nearly impossible to identify where mass-market software or hardware components might be deployed. As a result, we limit our analysis to the supplier's industry type.

Since our methodology allows for a more granular analysis to be performed as well, future studies can consider other categories, depending on the availability of reliable data.

Therefore, for our study we consider following factors as confounding variables:

1. public availability of the POC per vulnerability
2. availability of a fix – a mitigation or a patch
3. type of supplier' industry/sector
4. if supplier (primarily) delivers OSS or not

Matched case-control design applied in this study allows to change values of these parameters (e.g. go for different or more granular categories) or take into account additional confounding variables if reliable data about them is available. Selected confounding variables are considered when defining a similarity function to match cases and controls during their selection process.

Importantly, not considering other factors discussed above (for which there are no reliable data at this point) is mitigated through the randomisation element of the study design. The use of bootstrapping will minimise the impact of all possibly confounding or other external factors that we could not consider.

5.2.3 Similarity Function

Our independent and confounding variables for vulnerability v_i are listed in Table 5.

Variable	Description	Type	Values
Independent	Exposure to suppliers' PSIRT	Binary	g_i=1 if PSIRT g_i=0 otherwise
Dependent	'Exploited in the wild' status	Binary	e_i=1 if exploited e_i=0 otherwise
Confounding	Availability of POC	Binary	p_i=1 if POC available p_i=0 otherwise
Confounding	Availability of a fix	Binary	f_i=1 if fix available f_i=0 otherwise
Confounding	Open-source status of the supplier	Binary	o_i=1 if OSS o_i=0 otherwise
Confounding	Type of the industry of the supplier	Categorical	$t_i \in$ {software, hardware, cloud services, industry and manufacturing, cybersecurity industry}

Table 5: the overview of variables, types, and values they may take

Each vulnerability v_i can therefore be described with a vector of values:

$\mathbf{v_i} = [g_i, e_i, p_i, f_i, o_i, t_i]$,

A matched-case study requires us to match cases and controls according to the confounding variables p_i, f_i, o_i and t_i, to minimise their influence on the study results. We define similarity function between vectors of vulnerability v_i and v_j as

$$(3) \qquad s(\mathbf{v_i}, \mathbf{v_j}) = \begin{cases} 1 & \text{if } p_i=p_j \wedge f_i=f_j \wedge o_i=o_j \wedge t_i=t_j \wedge e_i \neq e_j \\ 0 & \text{otherwise} \end{cases}$$

In other words, similarity among two vulnerabilities with different status of 'exploited in the wild' (which is by design a selection of cases and controls we compare) is present if availability of POC, fix, type of industry, and OSS type of both vulnerability vectors are equal.

5.3 Datasets

A matched case-control design relies on robust, comprehensive data regarding disclosed and exploited vulnerabilities, ensuring that each vulnerability is represented by a complete vector of relevant parameters. Publicly available datasets curated by trusted authorities – continuously enriched by contributions from the global cybersecurity community – offer a valuable foundation for this purpose. In combination with additional databases and manually compiled supplier lists, these datasets provide the necessary basis to support the study. However, beyond outlining the used datasets, it is essential to consider alternatives to main data sources, assess their strengths and limitations, and address potential biases introduced by the chosen datasets.

5.3.1 National Vulnerability Database (NVD)

NVD[32] is the most authoritative, comprehensive, and trusted global, publicly accessible repository of disclosed vulnerabilities. It is maintained by the US National Institute of Standards and Technology (NIST). It contains records of about over 250,000 reported vulnerabilities, collected since 1999 (NIST, n.d.). Each record contains numerous relevant information, among other:

- a unique Common Vulnerabilities and Exposures (CVE) ID number
- a CVSS score
- references to advisories, mitigations, POC information, etc.
- an indication if a vulnerability is recorded in the catalogue of the known exploited vulnerabilities
- known affected software configurations and suppliers (each unequivocally labelled with an official Common Platform Enumeration (CPE) number)

The NVD aggregates data from various sources and serves as the largest publicly accessible repository for documented vulnerabilities. However, it often lacks the depth of technical detail provided by corporate databases maintained by major suppliers (such as Microsoft, Cisco, Oracle, and Apple), which offer more specific insights, including affected product versions, root-cause analyses, patching procedures, and specific mitigation steps. Nevertheless, these repositories predominantly focus on the supplier's own products or those directly related to their ecosystem, which introduces biases. Additionally, comprehensive vulnerability datasets maintained by leading cybersecurity firms (e.g. Rapid7, Kaspersky, or Palo Alto Networks) often require subscription-based access, creating barriers to broad, unbiased analysis. While these datasets are rich with contextual information, they may reflect the interests and threat models of specific clients.

Still, vulnerabilities disclosed through corporate as well as other channels are eventually registered with the NVD, as the assignment of a CVE ID is a prerequisite for broader recognition. This ensures that NVD maintains the most comprehensive dataset, albeit with less technical details per entry compared to corporate databases. The scope of vulnerabilities

[32] Available at: https://nvd.nist.gov/vuln

catalogued by NVD, alongside its role as a global aggregator, helps mitigate supplier bias, making it a reliable foundation for large-scale analysis. While the absence of broader technical details may limit the analysis of other possible risk factors, this study addresses such gaps by applying a robust methodology – including bootstrapping – to minimise potential confounding factors that are not of our interest. Future research of the relevance of specific factors for the likelihood of vulnerability exploitation could benefit from integrating multiple data sources to enhance the depth of information, capturing nuances that a single repository might overlook, and enhancing the applicability of this study.

5.3.2 CISA Known Exploited Vulnerability (KEV) List

KEV[33] is an authoritative public list of CVEs known to be exploited in the wild, maintained by the US Cybersecurity and Infrastructure Security Agency (CISA). KEV was established in November 2021 following the US Directive on 'Reducing the Significant Risk of Known Exploited Vulnerabilities' (CISA, 2021). It is intended to instruct US institutions and others about prioritising patching of certain critical vulnerabilities. KEV records all the vulnerabilities reported that are exploited in the wild, using the CVE number as identifier. Importantly, KEV lists only the vulnerabilities for which a patch has been published.

The list is not exhaustive, since there might be actively exploited vulnerabilities captured by other researchers that are not yet listed in the database. However, it remains the most comprehensive and trusted list publicly available. Other lists, maintained primarily by the cybersecurity industry, may be richer and more detailed, but are typically closed for subscribers only.

While criteria which drives CISA to include a CVEs to the KEV list are clear – an assigned CVE ID, evidence of exploitation in the wild, and availability of remediation (a patch or other fix) – it is expectable that their primary interest is to capture CVEs that endanger critical sectors in the USA, and thus might miss some of the other existing and exploited CVEs. This may introduce certain geographical bias. What is important for this research, however, is not if the list is exhaustive, but that the list is valid – what is in the database has certainly been exploited in the wild regardless of geography.

The KEV dataset could also be expanded by drawing on additional commercial or publicly available repositories of zero-day vulnerabilities – namely vulnerabilities exploited before public disclosure – such as Google's Project Zero[34] database or the CZ Zero Day[35] listings. A preliminary cross-comparison of these sources with the KEV catalogue at the time of this research, however, yields only approximately twenty additional CVEs, suggesting that the KEV dataset provides sufficiently broad coverage for the purposes of this study. Further sources could include investigative reports by cybersecurity companies, particularly those

[33] Available at: https://www.cisa.gov/known-exploited-vulnerabilities-catalog
[34] Available at: https://googleprojectzero.blogspot.com/
[35] Available at: https://www.zero-day.cz/database/

documenting APT campaigns relevant to state-level operations. Such material, however, is fragmented and often restricted by paywalls, limiting its suitability for systematic analysis. Importantly, the methodology developed here remains extensible and allows for future incorporation of additional public datasets as they become available.

5.3.3 Database of Suppliers

Each vulnerability is directly related to its supplier through the CPE identifier in the NVD record. Therefore, a list of tens of thousands of unique suppliers of vulnerable digital products could be exported from NVD. However, this list doesn't include attributes related to the operation of a PSIRT, open-source nature, or type of industry, which are needed for our analysis. There is no other comprehensive and reliable list with such information either.

Therefore, a required database of suppliers was manually aggregated for this study based on the publicly available information. Due to practical limitations, a limited number of suppliers of vulnerable products were selected to include representatives of all the industry types we defined for the study, both proprietary and OSS, that have visible capacities to provide a response of some sort at least (such as through issuing security advisories; or a PSIRT in the best case). The information about existing PSIRTs were collated from suppliers' websites and white papers, as well as from the list of members of various incident response and vulnerability management fora, like the PSIRT Special Interest Group[36] of the Forum of Incident Response and Security Teams (FIRST), and the assigners authorised by the CVE Program[37] to assign CVE IDs.

The database also includes a number of major OSS distributions and communities that possess the organisational capacity to operate PSIRT-like functions and to engage with bug-bounty programmes. Small-scale OSS projects maintained by limited or informal contributor groups were not included, as they generally lack the resources and structures necessary to implement such practices. For these projects, alternative mechanisms – such as platform-level security support (e.g. GitHub), umbrella OSS foundations, or community-based response models – may play a comparable role; however, these arrangements fall outside the scope of the present research.

Since our observational units are vulnerabilities (described by exploitation, POC, and fix), each related to a supplier, it is important that the sample of vulnerabilities from which we draw cases and controls is sufficiently big and diverse. This allows us to apply bootstrapping, which then compensates on the modest size of a supplier database, and minimise eventual bias due to its distribution. Nevertheless, it is still important to interpret the results in context of the chosen sample of suppliers. Importantly, methodology allows for a different or more comprehensive list of suppliers and their features to be fed into the study in future, as information becomes available.

[36] Available at: https://www.first.org/global/sigs/psirt/
[37] Available at: https://www.cve.org/partnerinformation/listofpartners

5.3.4 Exploit-db (EDB)

EBD[38] is a widely used and trusted public repository of known POCs for vulnerabilities, managed by a corporate sector. It maps all the CVEs for which there have been reported POCs, typically by the cybersecurity researchers. While some of the CVEs with POC might have also been exploited in the wild (though EDB doesn't report this), the existence of a POC doesn't imply the exploitation in the wild per se. This list is not exhaustive, but is considered authoritative and unique.

5.4 Data Pre-processing

Our observational unit is individual vulnerabilities, each represented by a CVE ID recorded in the NVD. For each vulnerability, additional parameters – such POC availability, exploitation status, and supplier information – were aggregated from other databases described in Section 5.3. Given the study focus on evaluating the efficiency of PSIRT, the dataset required specific pre-processing to ensure relevance and consistency.

The selection of vulnerabilities was limited to a specific time frame reflecting the period during which PSIRTs became a more standardised and widely adopted practice. While early iterations of product security teams began emerging among leading technology companies in the early 2000s (Rajnovic, 2010), wider implementation across industries gained traction following the updated publication of PSIRT Services Framework by FIRST in 2019 (FIRST, 2020b), as well as the introduction of PSIRT details within the international standards like ISO/IEC 30111 (ISO, 2019). These developments informed the decision to select 2020 as the starting point for the study, capturing vulnerabilities discovered after the establishment of these frameworks. Data points prior to 2020 were excluded to avoid addressing periods of inconsistent PSIRT deployment.

The observation period extends until July 2024, aligning with the conclusion of this research. Notably, in mid-2024 NVD experienced turbulations – processing delays and backlog, availability issues with API, CVE format transition, and the adoption of the CVSS4.0 (NVD News, 2024) – which would necessitate adjustments to the data processing workflow for newer entries. NVD and KEV databases were then filtered to preserve only vulnerabilities related to suppliers listed in the database of suppliers.

To create a possible pool of samples for choosing the matching controls, we then exclude all the CVEs contained in KEV from the formed NVD dataset. This ensures that the pool of possible controls does not include cases.

Importantly, data pre-processing does not impact the robustness of the analysis, since the resulting datasets contain sufficiently large numbers of entries – not least because the number of documented vulnerabilities continues to rise year-over-year (NIST, no date) – to preserve

[38] Available at: https://www.exploit-db.com/

methodological coherence without compromising data integrity. Number of vulnerabilities and corresponding unique vendors in each pre-processed database is shown in Table 6.

Set	No. of vulnerabilities	No. of suppliers
Possible cases (pre-processes KEV)	625	128
Possible controls (pre-processes NVD)	45325	339

Table 6: The number of vulnerabilities and suppliers in each pre-processed database

Pre-processed NVD database was finally filtered from vulnerabilities of vendors listed in the pre-processed KEV list, to ensure that the pool of possible controls doesn't contain any cases.

6 Model Deployment and Calculation

Having established the methodological foundation, we now turn to translating theoretical constructs and pre-processed data into actionable insights through statistical analysis. By implementing techniques like sampling, matching, and bootstrapping, we aim to ensure the validity of the results and isolate external factors. The chapter then introduces the risk reduction metrics, and explains the statistical tests employed to confirm the robustness of our findings, and to evaluate the relevance of confounding variables. Together, these calculations provide the empirical basis for understanding the efficiency of PSIRTs in reducing risks from vulnerability exploitation.

6.1 Sampling, Matching and Bootstrapping

By aggregating relevant information from pre-processed KEV, NVD, EDB, and our database of suppliers, we obtain:

- a set of possible cases:
  $C_p = \{\boldsymbol{v_1}, \boldsymbol{v_2}, ..., \boldsymbol{v_n}\}$, with $\boldsymbol{v_i} = [g_i, e_i, p_i, f_i, o_i, t_i]$, $\forall i \in \{1..n\}$, where n is a total number of vulnerabilities in a pre-processed KEV database
- a set of possible matching controls:
  $M_p = \{\boldsymbol{v_1}, \boldsymbol{v_2}, ..., \boldsymbol{v_m}\}$, with $\boldsymbol{v_j} = [g_j, e_j, p_j, f_j, o_j, t_j]$, $\forall j \in \{1..m\}$, where m is a total number of vulnerabilities from the pre-processed NVD database

Each vector $\boldsymbol{v}$ is linked to a unique CVE ID and the related supplier.

We then turn to (Figure 9):

1) sampling – creating a set of cases C, with one randomly selected vulnerability per each distinctive supplier
2) matching – creating a set of controls M where each control is matched with a single case
3) bootstrapping – repeating the process with randomisation

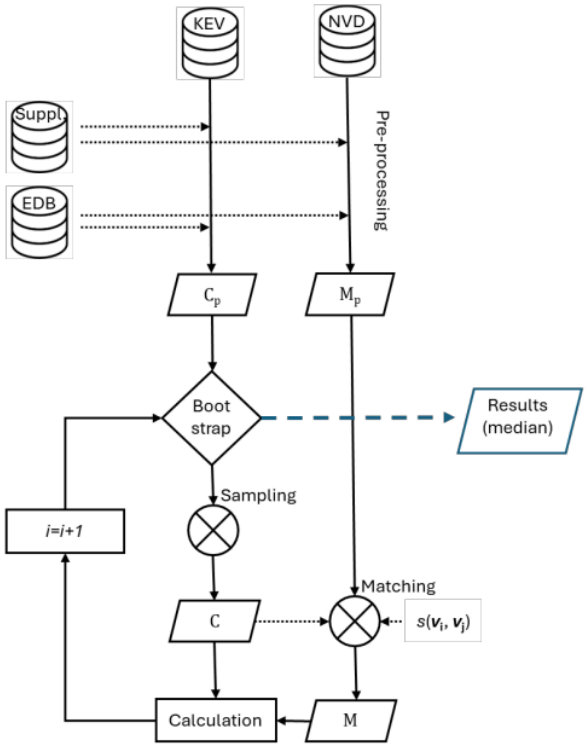


Figure 9: PSIRT efficiency assessment model based on matched case-control methodology with bootstrapping

6.1.1 Sampling

By its nature, some suppliers have multiple vulnerabilities in KEV list. This shouldn't come as surprise, since end-user platforms (like Microsoft, Apple or Google) are dominant targets of vulnerability exploitation (Stone et al., 2024). To address possible bias introduced by dominance of some suppliers, we normalise the database to contain only one entry 'representing' each distinctive supplier; the choice of a particular vulnerability to represent each supplier in KEV is made randomly. This involves following steps:

1. for each distinctive supplier featuring in C_p, we create a subset $C_{ps} \subset C_p$, consisted of all vulnerabilities of that supplier that are in C_p
2. from each C_{ps} we then select single v uniformly at random, to form set C of cases, which therefore contains a single representative vulnerability for each distinctive supplier

As part of the bootstrapping process (described below), multiple iterations of such random choice are conducted, reducing the effect of supplier's dominance.

6.1.2 Matching

For each case we now select a control with matching characteristics based on the similarity function in formula (1):

1. for each $\boldsymbol{v_i} \in C$ we create $M_{pi} \subset M_p$, by selecting all vectors $\boldsymbol{v_j}$ that satisfy the similarity function $s(\boldsymbol{v_i}, \boldsymbol{v_j})$
2. from each M_{pi} we then select single $\boldsymbol{v}$ uniformly at random, to create set M of matching controls

The number of elements in C and M is same by design: $|C| = |M|$, and C and M have the same distribution of confounding attributes.

6.1.3 Bootstrapping

Multiple vulnerabilities per supplier in C_p allow us to resample elements of C; similarly, large pool of possible matching controls in M_p enables us to resample elements of M. If we repeat resampling uniformly at random multiple times – a process known as bootstrapping – we reduce the bias of dominant suppliers as well as the impact of external and unobserved factors. Bootstrapping is particularly useful in such cases as it does not rely on assumptions about the underlying distribution of data.

6.2 Risk Reduction

Our presumption is that presence of a PSIRT of a supplier reduces the likelihood of exploitation of its vulnerabilities in the wild. To test our presumption and look for how big the effect of a practice on likelihood is, we turn to calculating risk reduction.

We compute frequency with which exposure to a risk factor – in our case, good practice PSIRT – identifies a vulnerability exploited in the wild. We calculate values of *a*, *b*, *c,* and *d*, as defined in Table 4, where G denotes a set of all vulnerabilities $\boldsymbol{v_i}$ with value $g_i = 1$ (i.e. PSIRT is in place).

Firstly, we obtain:

$a = |\{\boldsymbol{v} \in C \mid G\}|$

$b = |\{\boldsymbol{v} \in M \mid G\}|$

$c = |\{\boldsymbol{v} \in C \mid \neg G\}|$

$d = |\{\boldsymbol{v} \in M \mid \neg G\}|$

where *G* is an attribute marking the exposure to a good practice.

Based on this, we calculate risk reduction and odds ratio.

Experimental Event Rate (EER) represents a risk of an outcome related to exposed samples, i.e. the rate of occurrence of the exploitation among samples exposed to the PSIRT. It is a

probability of the outcome (exploitation) among vulnerabilities whose suppliers are exposed to a good practice:

$$EER = P(v \in C \mid v \in G) = P(v \in C \wedge v \in G) / P(v \in G)$$

EER is, therefore, calculated as a total number of cases exposed to a good practice, divided by a total number of cases and samples exposed to a good practice. Therefore:

(4) $$EER = a/a+b$$

Control Event Rate (CER) represents a risk of an outcome related to unexposed samples, i.e. the rate of occurrence of the exploitation among samples that were not exposed to the good practice. It is a probability of the outcome (exploitation) among vulnerabilities whose suppliers are not exposed to a good practice:

$$CER = P(v \in C \mid v \notin G) = P(v \in C \wedge v \notin G) / P(v \notin G)$$

CER is, therefore, calculated as a total number of cases not exposed to a good practice, divided by a total number of cases and samples not exposed to a good practice. Therefore:

(5) $$CER = c/c+d$$

Risk Reduction (RR) measures the reduction in risk associated with a preventive measure, i.e. the strength of association between exposure to PSIRT as a good practice, and the risk of vulnerability exploitation in the wild:

$$RR = P(v \in C \mid v \in G) / P(v \in C \mid v \notin G)$$

Therefore:

(6) $$RR = EER/CER$$

Relative Risk Reduction (RRR) tells us in absolute terms how much lower the modified risk of the outcome is from the start risk:

$$RRR = 1 - RR$$

Therefore:

(7) $$RRR = (CER-EER)/CER$$

Absolute risk reduction (ARR) measures the extent to which a preventive practice reduces the risk of an outcome i.e. exploitation in the wild. ARR is a difference in probability of the exploitation among exposed and non-exposed samples:

$ARR = P(v \in C \mid v \notin G) - P(v \in C \mid v \in G)$

Therefore:

(8) $ARR = CER - EER$

ARR is the most useful parameter for our study, as it directly assesses the efficiency of a good practice: it quantifies the extent to which presence of PSIRT reduces the occurrence of exploitation. ARR allows us to, regardless of the baseline risk, interpret the absolute magnitude of risk reduction due to a good practice.

Odds ratio (OR) indicates the strength of the association between an exposure and an outcome. It compares odds of an exposure among cases to odds of exposure in controls:

(9) $OR = (a/c) / (b/d)$

OR tells us if it is more likely that a supplier of a vulnerability exploited in the wild has PSIRT than if it doesn't.

6.3 Testing Statistical Significance

Besides calculating risk reduction, we must also test whether the independent variable has a statistically significant relationship with the dependent variable. For this, we firstly define our null hypothesis:

> H_0: exploitation of a vulnerability and the good practice in place are not related in the population; the proportion of vulnerabilities exploited is the same regardless of the presence of a good practice.

In other words, we assume a null hypothesis of no relationship between the two variables.

Our alternative hypothesis is:

> H_a: exploitation of a vulnerability and the good practice in place are related in the population; the proportion of vulnerabilities exploited is not the same in case of absence or presence of a good practice.

Therefore, our alternative hypothesis is that there is a significant association between the two variables.

If analysis shows that there is significant association between the two variables, we can reject the null hypothesis and support the alternative one. Otherwise, we fail to reject the null hypothesis, and our alternative hypothesis does not hold.

Our study is defined in the following way:

- our hypothesis is one-tailed (good practice decreases likelihood of exploitation)
- it has categorical (binary) outcome variable (exploited or not)
- it has categorical (binary) predictor variable (exposed to a good practice or not)
- randomised selection of controls to each sample

Accordingly, a non-parametric test, such as a *chi-square (χ^2) test of independence*, is appropriate to test our hypothesis, i.e. the association between the two variables. In calculating chi-square, we take that the ration of controls per case is r=1, since only one control is matched to each case. If the test statistic is larger than a critical value of a chi-square distribution, we will be able to reject the null hypothesis.

We also calculate the statistical significance, *p value*, which determines the likelihood that the study observation is due to chance. *Statistical significance level (α)* is the largest acceptable risk for a false positive conclusion; we set a common value of α=0.05, meaning that results must have below 5% chance of occurring in case of a null hypothesis. If the $p<\alpha$, we can conclude that observations are statistically significant, and we can reject the null hypothesis.

Statistical power ($1-\beta$) informs about sensitivity in terms of the likelihood of a significance test to correctly detect an effect. In our case, statistical power is calculated as the proportion of the chi-square statistics that exceeds the critical value of 3.841 (for α=0.05). Standard deviation of the outcome variable (σ) indicates the variability in the outcome within the population.

Confidence Intervals (CI) indicates the probable range of values for the true effect in population. CI will be calculated around χ^2, OR, and p-value. We use a commonly used 95% confidence interval. Interpreting the CIs is important for assessing the precision and certainty of the study.

7 Results

Building on the deployed model and methodology, this chapter presents the key findings that quantify the impact of PSIRTs on vulnerability exploitation risks. By analysing metrics such as RR, RRR, ARR, and OR, alongside statistical significance indicators like chi-square tests and p-values, we aim to evaluate the reliability and practical implications of these results. These insights are critical for understanding the effectiveness of PSIRTs, setting the stage for a broader discussion on their significance and application.

7.1 Risk Reduction

We have run the entire experiment with bootstrapping for 200 times, above which results tend to concentrate around the same value (Figure 10). All the statistics reported below are the median of the generated distribution of values.

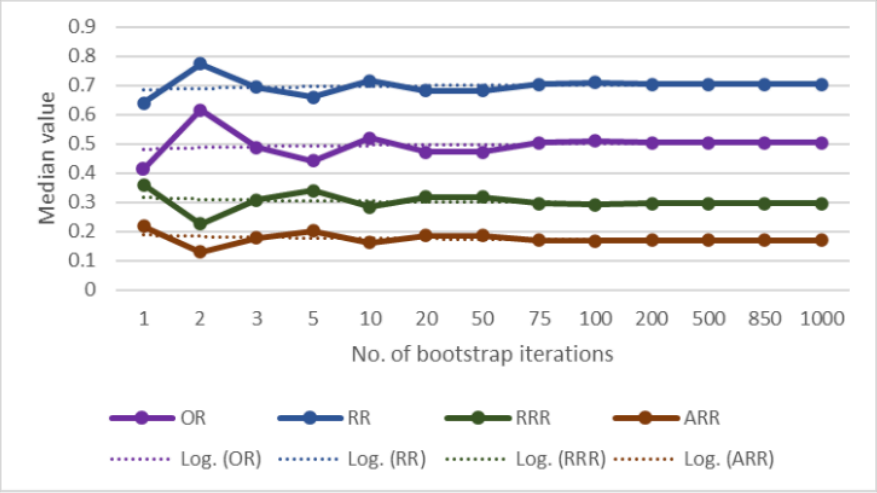


Figure 10: Median values of risk reduction parameters depending on number of bootstrap iterations

Results related to risk reduction are shown in Table 7.

Parameter	Value	Percentile
Risk reduction (RR)	0.7044	70%
Relative risk reduction (RRR)	0.2956	30%
Absolute risk reduction (ARR)	0.1701	17%
Odds ratio (OR)	0.5028	50%

Table 7: Median results related to risk reduction, obtained with 200 iterations of the experiment

2

RR<1 tells us that vulnerabilities of suppliers with PSIRT have a lower risk (70%) of the exploitation in the wild compared to those without it. This result suggests that presence of PSIRT may be associated with a reduced risk from exploitation.

RRR=0.3 confirms a proportional risk reduction compared to the baseline risk, yet a relatively low RRR signals a modest reduction. This means that PSIRT, as a good practice, discerns between high and low risks of vulnerability exploitation at a moderate level.

ARR=0.17 tells us that the PSIRT as a practice is associated with 17% decrease in the risk of exploitation of a vulnerability, compared to not having this preventive practice in place. Since ARR directly assesses the efficiency of a good practice, having in mind RRR as well, we can interpret this result as a relatively modest efficiency of PSIRT as a good practice.

OR=0.5 indicates that vulnerabilities by suppliers with PSIRT have 0.5 times the odds of being exploited in the wild than those of suppliers without PSIRT. With a 95% confidence interval (CI) of (0.36, 0.67) for OR, we are 95% confident that the OR in the population would be between 0.36 and 0.67. This CI is somewhat wider, suggesting OR variability. However, values deviate approximately ±30% of the midpoint OR, and the CI range spans less than one order of magnitude, which can be considered relatively narrow for a tight 95% in observational studies with real-world data and OR variability. In context of our study design and research question, the CI has enough precision to support the usefulness of the findings about the reduced odds of vulnerability exploitation for suppliers with PSIRT. Additional precision may be achieved in future if richer public databases of PSIRTs become available.

In Figure 11, we also show results when only certain confounding variables are taken into consideration: *p*, *o* or *t* (by, design, *f* has to be considered, since vulnerabilities in KEV database are characterised by *f*=1). This allows us to re-assess their impact on the likelihood. From those results, we see that the exclusion of POC as a confounding variable increases the ARR by over 10%, which confirms that POC impacts the likelihood significantly, and that we correctly considered it as confounding. Type of industry and OSS, on the other hand, increase ARR by only 3.6% and 2.2% respectively, signalling that they do impact the likelihood, but not significantly.

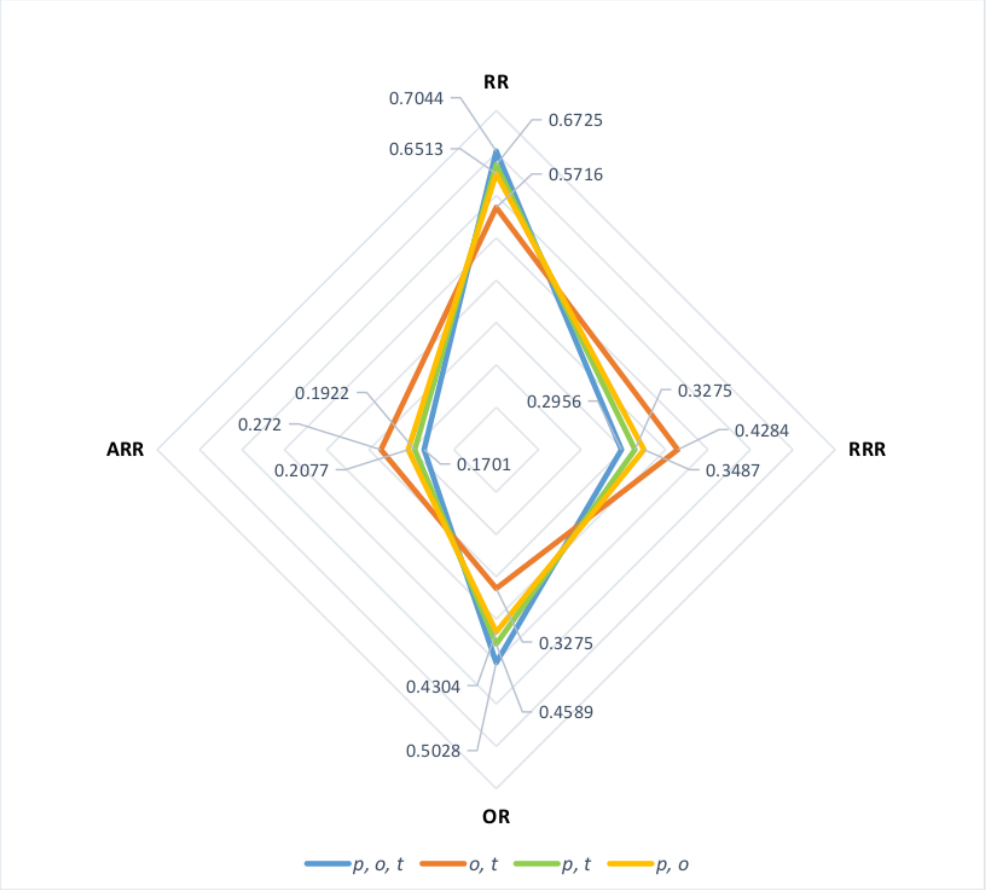


Figure 11. Median results related to risk reduction, depending on which confounding variables are considered (p – POC, o – OSS, t – type of industry)

7.2 Statistical Significance

Results related to statistical significance of our model are shown in Table 8.

Parameter	*Value*
χ^2 test statistics	7.133152174
χ^2 test of independence	$\chi^2 > 3.841$
χ^2 confidence intervals (95%)	(2.00, 14.41)
p-value (<0.05)	0.003783579
confidence intervals (95%)	(0.00007-0.078)
Level of significance (α)	0.05
Statistical power (1 – β)	0.817
Standard deviation (σ)	3.24

Table 8: Median results related to statistical significance, obtained with 500 iterations of the experiment

With a degree of freedom df=1 and α =0.5, critical value of chi-square for a single tail is 3.841. Median χ^2 = 7.13, thus χ^2 > 3.841, signals the success of the test of independence. Value of χ^2 is moderately high, indicating a reliable association between good practice exposure and exploitation in the wild. Further, the confidence interval (2.00, 14.41) for 95% confidence doesn't include low values, indicating a consistent and strong association of the variables across the bootstrapped samples. Moreover, median p = 0.0038 and α=0.05, thus $p<\alpha$, suggests that the observed association is statistically significant, while a low confidence interval for the p-value indicates that this result is not due to random chance.

This leads us to conclude that there is strong evidence to suggest a statistically significant association between PSIRT and exploitation of vulnerabilities. This allows us to reject the null hypothesis H_0 and support the alternative hypothesis H_a of significant association between the two variables.

Statistical power of 0.82 indicates that there is 82% chance of detecting a true effect. Since it is above a statistical power of 0.8, which is considered adequate for most studies, we can conclude that the study design is effective at detecting a meaningful association between the variables. Standard deviation σ=3.24 suggest variability in the chi-square values, but given the overall results, it does not raise significant concerns.

8 Discussion

This discussion chapter interprets the empirical findings of the study in light of the theoretical, technical, and policy-oriented frameworks developed in the preceding chapters. Having established the significance of vulnerabilities in digital products for the execution of sophisticated OCO, mapped existing industry and academic practices for vulnerability reduction, and tested a novel model for assessing their effectiveness on a PSIRT case-study, the discussion now turns to the broader implications of those results. In particular, it reflects on what the findings reveal about the current state of industry practices, the role and limitations of PSIRTs, and the extent to which security-by-design approaches can meaningfully affect real-world exploitation dynamics.

The chapter proceeds by critically examining the maturity and coverage of industry practices across the digital product lifecycle, highlighting both areas of consolidation among leading suppliers and persistent gaps across regions, sectors, and organisational sizes. It then discusses the empirical results related to PSIRT effectiveness, including practical significance, cost considerations, and the influence of external factors such as POC availability and industry type. Finally, the discussion puts these findings within a broader geopolitical and strategic context, assessing how improved vulnerability management may increase the costs, uncertainty, and risk associated with offensive cyber operations. In doing so, it links technical practices at the organisational level to their potential systemic effects on cyber conflict and cyber armament.

8.1 A Mature but Uneven Landscape of Industry Practices

The analysis of academic literature, standards, and industry-led frameworks, conducted in Chapter 3, suggests that leading suppliers of digital products have indeed developed many recognised good practices for secure design and vulnerability management across the product lifecycle. Threat modelling, code analysis, TPC dependency management (such as SBOM), and monitoring and management of vulnerabilities are now well-established elements related to the SDL approach within major technology firms. These practices are reinforced by international standards (ISO/IEC), sector-specific frameworks (e.g. NIST), and voluntary alliances such as SAFECode and OWASP. In this sense, security-by-design has moved beyond a conceptual ideal and become embedded in the development processes of many technology vendors. Academic models and industry frameworks largely converge when it comes to security-by-design approaches throughout the lifecycle, indicating a degree of common understanding around what constitutes ‘good practice’.

At the same time, the evidence highlights significant differences in both the depth and consistency of implementation. While security principles are often integrated at early stages of development, their application is frequently constrained by market pressures, legacy architectures, and the increasing complexity of modern software ecosystems. In particular, product development heavily relies on TPC – especially the OSS – while the lack of transparency in complex global supply chains continue to transfer vulnerabilities that are not

fully mitigated by existing SDL practices. Moreover, empirical studies and industry maturity models show that proactive risk reduction (i.e. preventing the emergence of vulnerabilities in the first place) still lags behind reactive practices such as vulnerability scanning and patching. This imbalance suggests that, despite frequent commitments to 'shifting left' and introducing security from the very start of product development, organisations still prioritise investing in detection and remediation over preventive design choices – especially where preventive measures require organisational change, longer development timelines, or higher costs.

A further limitation concerns measurement and verification. While numerous frameworks assess process maturity or practice adoption, relatively few link these metrics to real-world security outcomes – particularly to the exploitation of vulnerabilities in the wild, as an essential component of sophisticated attacks. Industry models such as BSIMM and SAMM provide valuable insight into how widely specific practices are adopted, but they don't establish whether those practices measurably reduce exploitation. Academic approaches also tend to focus on efficiency indicators (e.g. patching speed or detection rates) rather than on exploitation risk itself. As a result, although it is reasonable to conclude that the recognised practices that leading suppliers have adopted are indeed good, it remains less clear which practices meaningfully reduce the strategic value of vulnerabilities, especially in the context of sophisticated, state-linked OCO.

Within this broader landscape, the emergence of PSIRTs represents a particularly important organisational development. PSIRTs operationalise vulnerability management by centralising responsibility for receiving reports about vulnerabilities and their assessment, coordination, public disclosure, and a release of a fix, as well as by connecting internal development teams with external researchers, CERTs, and customers. Moreover, PSIRTs contribute to developing key elements of the 'shift left' approach within organisations. Among leading suppliers, PSIRTs are increasingly formalised through documented disclosure policies, defined coordination and escalation paths, and participation in international cooperation mechanisms. Academic and practitioner sources broadly agree that PSIRTs enhance coherence, transparency, and responsiveness in vulnerability handling, especially in complex environments involving TPC and cross-border and cross-stakeholder coordination. The widespread endorsement of PSIRTs by standards bodies (ISO/IEC), coordination communities (e.g. FIRST), and large vendors, suggests that PSIRT has become a de facto good practice for mature vulnerability management.

Nevertheless, the discussion also reveals that establishing and maintaining a PSIRT requires substantial organisational and financial resources and therefore remains mainly the domain of larger software companies – typically those with more than 1,000 employees –predominantly based in developed countries or states with strategic investments in technology. Their strong presence among large multinational firms contrasts with their limited adoption by smaller suppliers and many open-source projects. Furthermore, the possibility to engage a PSIRT proactively rather than just reactively – thus to shape secure design and dependency governance – depends on their authority and mandate, resources, integration with other teams, and overall organisational culture.

Importantly, while PSIRTs clearly improve coordination and remediation processes, their contribution to reducing exploitation risk must be assessed in relation to external conditions – such as exploit availability, attacker incentives, and target attractiveness – which lie partly beyond organisational control. These observations reinforce the need for empirical approaches that move beyond maturity claims and directly examine outcomes, setting the stage for the methodological contribution developed in the previous sections.

8.2 Association Between PSIRT Adoption and Lower Exploitation Risk

Based on the theory and practice reviewed, PSIRT is expected to contribute measurably to reducing the likelihood of vulnerability exploitation in the wild. Since literature lacks a methodology to quantify the PSIRT impact on active vulnerability exploitation, the deployment of our methodology offers some useful insights.

Our findings confirm that the establishment and operation of a PSIRT contributes to measurably reducing the likelihood that disclosed vulnerabilities will be exploited in the wild. All four risk-reduction parameters considered – RR, RRR, ARR, and OR – consistently point to a statistically significant reduction in the likelihood of vulnerability exploitation among suppliers that maintain a PSIRT. Tests of statistical significance indicate that the model performs robustly in evaluating the effectiveness of an industry practice aimed at reducing exploitation in the wild.

8.2.1 Significance and Practical Impact of PSIRT

The empirical results demonstrate that the presence of a PSIRT is associated with a statistically significant reduction in the likelihood of vulnerability exploitation in the wild, with an observed ARR of 17%. In practical terms, this means that for every 100 disclosed vulnerabilities, approximately 17 fewer are expected to be actively exploited when a supplier maintains a PSIRT. As ARR directly reflects the effectiveness of a preventive practice, this finding provides strong empirical support for treating PSIRT as a meaningful industry good practice in vulnerability management.

At the same time, the magnitude of the effect must be interpreted in light of the organisational and financial costs of establishing and maintaining a PSIRT. Dedicated staffing, specialised training, international cooperation, and close integration with development, legal, and communications processes require significant investments. From an industry perspective, this raises legitimate questions about cost-effectiveness, particularly for smaller suppliers or those operating with constrained security budgets. A 17% reduction may appear modest when assessed in isolation.

However, in context of real-world developments, the practical value of this reduction becomes clearer. For suppliers managing large number of vulnerabilities in their products – such as

major software vendors or cloud service providers – even a partial reduction in exploitation likelihood can translate into substantial risk mitigation when at scale. In cloud environments, this may reduce losses associated with data integrity breaches, computational damage, or cascading attacks across neighbourhoods (Zhang et al., 2014). In industrial control systems (ICS) and operational technology (OT) – especially in context of critical infrastructure, like powerplants – where a single exploited vulnerability in embedded devices, programmable logic controllers (PLCs), or SCADA systems can trigger serious economic damage or physical harm, even relatively modest reductions in attacks probability can yield disproportionately high benefits (Parsons, 2022; Ravindrababu and Alves-Foss, 2022).

The interpretation of these findings must also be situated within the broader landscape of security-by-design and vulnerability management practices. We should, therefore, compare the efficiency of PSIRTs with other recognised practices. As discussed in Section 4.3, existing literature identifies numerous SDL metrics, yet different indicators are applied in various, often narrow, case studies (Kudriavtseva and Gadyatskaya, 2024). Most proposed metrics assess process maturity or operational efficiency rather than real-world exploitation outcomes, and few are empirically validated or widely adopted by practitioners. To the best of the author's knowledge, no prior study has quantified the effectiveness of a vulnerability management practice in terms of reducing active exploitation, which makes direct comparisons difficult. The proposed methodology will allow this comparison in future work.

Reliability of data about PSIRT practices should also be considered. Although reviewed literature provides many related standards, guidelines and examples, there is no standardised PSIRT model, and suppliers tend to follow their own needs, experience, intuition, and resources. Publicly available information about PSIRTs – typically collected from vendor websites or membership in coordination communities – may therefore be imperfect or incomplete. Some suppliers may formally claim to operate a PSIRT while providing limited capacity, whereas others may perform PSIRT-equivalent functions without explicitly labelling them as such. Greater standardisation of PSIRT practices, along with improved public mapping, could strengthen future assessments of effectiveness.

Finally, PSIRTs rarely operate in isolation. In practice, their presence often signals broader organisational commitment to secure development, including threat modelling, supply-chain and TPC management, secure deployment, and participation in bug-bounty or CVD ecosystems. These practices are likely mutually reinforcing, and their combined effect on reducing exploitation risks may exceed the impact of any single measure. While this study isolates PSIRT to enable empirical testing, further research is needed to separate and compare the contributions of individual practices and to explore their cumulative impact on reducing incentives and opportunities for OCO.

8.2.2 Impact of the POC and Industry Characteristics

The results of our simulation also demonstrate that the availability of POC is an important determinant of vulnerability exploitation in the wild, in line with earlier empirical work discussed in Sections 2.1.3.2 and 4.3.5. When POC availability is incorporated as a confounding factor in the model, the ARR associated with the presence of a PSIRT changes by 10%. This finding reinforces the importance of suppliers maintaining a dedicated capability – such as a PSIRT – to continuously monitor researcher publications, code repositories, social media, and underground fora for available POC and exploits, rather than focusing solely on vulnerability disclosures and patches.

By contrast, industry type and the open-source nature of products exhibit a more limited influence in this study. Accounting for industry sector alters the ARR by only 3.6%, while the open-source versus proprietary distinction accounts for a further 2.2%. This runs counter to common assumptions that certain sectors or open-source components are inherently more exposed. One possible explanation is the increasing trend of indiscriminate targeting and exploitation, particularly within cybercrime ecosystems, where widely available off-the-shelf exploits – such as those integrated into ransomware or automated scanning campaigns – are deployed widely, reducing the importance of the specific nature of the target.

At the same time, the relatively small effect attributed to OSS should be interpreted with caution. The dataset includes only a limited number of large open-source projects that possess sufficient organisational capacity to operate PSIRT-like structures, potentially underrepresenting the broader open-source ecosystem. A dedicated follow-up analysis focused specifically on open-source vulnerability management models, drawing on a dataset built for that specific purpose, would be required to draw more robust conclusions in this area.

Finally, additional granularity is needed to generalise these findings across organisational and sectoral contexts. The efficiency of PSIRTs is likely to vary with organisational size, available resources, and cybersecurity maturity. Larger suppliers may achieve greater risk reduction through more formalised processes, greater human resources, specialised tooling, and stronger integration with development pipelines, while smaller organisations may struggle to achieve similar benefits. Likewise, sectors subject to stringent regulatory regimes – such as finance or healthcare – have different baseline risk profiles and response capabilities due to higher compliance. Future research should therefore disaggregate PSIRT effectiveness across organisational scales and regulatory environments to refine understanding of where and how this practice delivers the greatest value.

8.2 Impact of PSIRT on Offensive Cyber Operations

The literature reviewed in Section 2 establishes vulnerability exploitation as a foundational enabler of cyber weapons and OCO. Sophisticated, politically motivated attacks – often conducted by the APT groups operating as state proxies – require sustained access to

ploitable vulnerabilities within specific products, services, or supply chains. For such actors, success depends not merely on identifying vulnerabilities, but on maintaining a window of opportunity long enough to operationalise, deploy, and repeatedly exploit them before mitigation measures are implemented on targeted systems. Their major interest, therefore, is to increase the likelihood of exploitation of a particular vulnerability in a specific system.

The empirical findings of this research indicate that the presence of a PSIRT reduces the likelihood of vulnerability exploitation in the wild by approximately 17%. Translated into operational terms, this reduction represents a meaningful decrease in the probability that an attacker will be able to successfully exploit a given vulnerability in real-world conditions. For widely deployed software platforms, cloud services, or digital components embedded in critical infrastructure, even a modest reduction in exploitation probability can significantly complicate adversarial planning and execution. Moreover, as discussed previously, PSIRT adoption rarely occurs in isolation: it is typically associated with broader investments in SDL, supply-chain risk and TPC management, coordinated disclosure processes, and organisational security awareness. Taken together, these layered practices further increase the cost, complexity, and uncertainty of mounting successful OCO.

Revisiting the specific hypothesis formulated in Section 1.4 – that the feasibility of OCO can be reduced through PSIRT impact on the exploitation likelihood – the results provide qualified support. The hypothesis holds most clearly within segments of the global digital supply chain dominated by large suppliers with mature security practices and sustained engagement in CVD. In these contexts, PSIRTs contribute to narrowing exploit windows, limiting repeatable exploitation, and raising the operational costs for attackers.

However, the analysis also reveals important limitations. PSIRT adoption remains unevenly distributed across geographies, industries, and organisational sizes. As shown in Section 3.4.2.3, most mature PSIRTs are concentrated within countries that possess strong digital industries and strategic interests in cyberspace, including the United States, EU member states, China, Japan, the Republic of Korea, Russia, and India. These are also jurisdictions that tend to invest heavily in cyber defence and offence. By contrast, smaller suppliers, firms operating in less digitally advanced economies, and many actors embedded deeper within global supply chains often lack comparable capabilities.

These gaps create persistent weak points in the supply chain that can still be exploited by state and non-state actors alike. Effective practices at large vendors may force adversaries to target weaker downstream suppliers and integrators instead. Consequently, while PSIRTs can meaningfully constrain OCO, their cumulative impact depends on the ecosystem-level rather than firm-level security, and thus a wider diffusion of secure-by-design practices across suppliers, sectors, and regions.

8.3 Impact of Industry Security Practices on Offensive Cyber Operations

The findings of this study invite a broader reflection on what effective vulnerability management means from the perspective of an attacker. For actors conducting sophisticated OCO, particularly state-backed groups and APTs, the failure of an operation due to pre-emptive mitigation by a supplier's security practices – such as those coordinated by a PSIRT – represents more than a missed opportunity. It entails the loss (i.e. 'burning') of a potentially valuable vulnerability, compressed window of opportunity for action, increased operational uncertainty, and heightened strategic risk. For such actors, failed exploitation attempts may lead not only to wasted technical investment but also to political, reputational, or diplomatic costs, especially where plausible deniability is challenged. In this sense, even modest reduction in exploitation likelihood may have disproportionate deterrence effects, forcing attackers to invest in a broader stockpile of vulnerabilities and exploits, or in shifting to alternative yet less favourable attack vectors (e.g. social engineering).

At the same time, the capacity of vulnerability management practices to disrupt OCO has inherent limits. Patching processes remain imperfect, particularly in operationally sensitive environments. Even when a vulnerability is addressed by a supplier, opportunities for exploitation may persist in other, unpatched deployments. This challenge is especially visible in critical infrastructure sectors. ICS and OT environments are complex, often rely on supply chains with multiple layers, and integrate numerous TPCs whose ultimate use cases (e.g. in critical operations) may not be visible to upstream suppliers. Moreover, patching in such environments is operationally costly: downtime, safety validation, and system re-integration create strong disincentives for rapid remediation. At the same time, critical infrastructure remains a high-value target for hybrid operations, where OCO can generate far-reaching economic or societal disruption with high strategic payoff and limited attribution.

Encouragingly, leading suppliers serving critical sectors increasingly adopt security-by-design approaches, including mature SDL frameworks and PSIRT functions, as discussed in 3.4. At the same time, regulatory frameworks have begun to reflect heightened geopolitical awareness of critical infrastructure vulnerabilities and the related national and international security concerns, imposing stricter security and supply-chain obligations. These efforts are reinforced at the international level through multilateral and regional processes – such as the UN OEWG and its successor a permanent UN Global Mechanism on Developments in the Field of ICTs in the Context of International Security and Advancing Responsible State Behaviour, cooperation and confidence-building measures within regional organisations like the OSCE, OAS and ASEAN, and multistakeholder and capacity building initiatives – which collectively increase awareness and provide policy guidance for both regulators and the industry (Kazakova et al., 2025).

However, increased geopolitical tensions and fragmented regulatory approaches risk creating new avenues for exploitation rather than closing existing ones. States increasingly consider requiring researchers to disclose vulnerabilities to national authorities rather than directly to vendors. While national CERTs can act as effective intermediaries, this arrangement becomes

problematic in environments where CERTs are closely subordinated to national security agendas, and where transparency of how vulnerabilities will be further dealt with is limited. In such cases, disclosed yet unpatched vulnerabilities may be deliberately stockpiled for intelligence or OCC, undermining broader societal security and weakening trust in CVD processes (Kazakova et al., 2025). This dynamic illustrates how vulnerability management can become entangled with cyber armament rather than serving as its constraint. On the other hand, by 'burning' the discovered vulnerabilities more efficiently, good practices may potentially accelerate volatility in cyber weapons stockpiles.

An additional structural factor shaping the longer-term impact of vulnerability management on OCO is the growing reliance on vulnerable OSS. The expansion of OSS ecosystems – accelerated by AI-assisted coding and the proliferation of open-source approach due to strategic benefits of the open AI models – will further embed OSS deeply within digital products and supply chains. Yet many OSS projects are developed and published by small, volunteer-driven communities with limited resources for systematic security practices. For instance, while most OSS projects nominally support CVD, empirical evidence suggests that only a fraction of reported vulnerabilities follow CVD timelines in practice (Liu et al., 2025). At the same time, attackers increasingly target supply chains through OSS, exploiting trust relationships or manipulating development processes, as demonstrated by incidents such as Log4j and XZ Utils (Hamer et al., 2025; Przymus and Durieux, 2025).

Some mitigating mechanisms are emerging. Large OSS distributions, repository platforms such as GitHub, and broader ecosystem actors increasingly provide tooling for vulnerability discovery and analysis, alerts, and automated remediation support (Kazakova et al., 2023). Academic and commercial research is also exploring the application of AI to vulnerability detection and patch generation in OSS (Pleskonjić et al., 2025). Importantly, recent policy and research discussions converge on the view that responsibility for OSS security should shift away from individual maintainers toward the industry that integrates, commercialises, and deploys OSS within products and services (Kazakova et al., 2023). In this context, industry practices such as SBOM implementation, TPC governance, and PSIRT coordination can generate positive spill-over effects by strengthening one of the weakest links in the global digital ecosystem. Consequently, this could significantly increase the costs on the OCO side.

Taken together, these observations allow a nuanced assessment of the general hypothesis (Section 1.4): that cyber armament and politically motivated OCO can be constrained through the systematic adoption of security-by-design and vulnerability management practices. The empirical results and broader discussion support this hypothesis, but only conditionally. PSIRTs and related practices demonstrably reduce the likelihood of vulnerability exploitation, yet their cumulative effect depends on wider adoption across sectors, geographies, and organisational sizes; alignment with regulatory and policy frameworks; and governance arrangements that explicitly address OSS and supply-chain dependencies.

Several implications follow:

1. Continued empirical assessment of the measurable effectiveness of security practices – using the proposed methodology with advanced datasets and real-world attack contexts (e.g. APT operations) – is essential to guide investment and policy decisions.
2. Broader uptake of security practices, including PSIRT maturity stages, among smaller suppliers and across diverse regions is necessary for system-level impact; leading firms with SDL experience can play a catalytic role through supply-chain requirements and multistakeholder cooperation platforms like FIRST or the Geneva Dialogue.
3. Companies that benefit from OSS must assume greater responsibility for monitoring and mitigating vulnerabilities within those components. Besides, regulatory and standardisation efforts should emphasise responsibility of industries that integrate and deploy OSS components to the market – and incentivise them to apply practices like SBOM.
4. Standardisation bodies and regulators should work on establishing interoperable standards related to security practices like SBOM or PSIRT, while avoiding jurisdictional fragmentation.
5. Finally, global and regional multilateral and multistakeholder processes remain crucial for shaping shared expectations in form of principles and norms, thus enhancing the commitment of signatory and supporting states and actors, and consequently incentivising good practices beyond the group of technologically advanced states and firms.

Where political will and economic incentives align, pathways toward more secure digital products clearly exist. Yet a persistent risk looms in the form of legal, technical, and economic fragmentation of global technology supply chains, driven by geopolitical polarisation and coupled with advances in AI and quantum technologies. Such developments could alter both the nature of OCC and the effectiveness of existing security practices, underscoring the need for continued interdisciplinary research at the intersection of technology, security, and international politics.

9 Conclusion

This dissertation addressed an important challenge in contemporary cybersecurity and international security studies: how vulnerabilities in digital products enable offensive cyber operations (OCO), and whether industry practices can meaningfully constrain this dynamic. Against the backdrop of intensifying cyber armament and the strategic stockpiling and exploitation of vulnerabilities by states and their proxies, the research examines whether security-by-design and vulnerability management practices adopted by suppliers of digital products can reduce the likelihood of exploitation in the wild (i.e. in real-world operations) and thereby increase the cost and uncertainty of OCO.

Existing literature and industry experience suggest that certain supplier-level security-by-design and vulnerability management practices can be identified as 'good practices' for reducing risks associated with vulnerability exploitation. Among them, the establishment of a Product Security Incident Response Team (PSIRT) stands out as a widely recommended organisational practice for managing vulnerability discovery, assessment, disclosure, and remediation. However, prior to this research, there was no dedicated empirical methodology to verify the effectiveness of PSIRT or similar practices in reducing real-world exploitation of vulnerabilities, nor to link such effects to broader dynamics of cyber armament and OCO.

This dissertation responded to that gap by combining:

- a multidisciplinary review of cybersecurity risks, vulnerability exploitation, and offensive cyber operations
- a structured analysis and classification of academic models, international standards, regulatory frameworks, and leading industry practices for secure development and vulnerability management
- the development and empirical deployment of a novel methodological model to assess the effectiveness of PSIRT as a good practice

The proposed model draws on a randomised matched case–control design, adapted from epidemiology and applied to cybersecurity data. It operationalises vulnerability exploitation risk using authoritative public datasets on disclosed vulnerabilities, known exploited vulnerabilities, proof-of-concept (POC) availability, and supplier characteristics, including the presence or absence of a PSIRT.

The principal findings of the study can be summarised as follows:

- *PSIRT impact on likelihood of exploitation*: The empirical results demonstrate that the presence of a PSIRT is associated with a statistically significant reduction in the likelihood of vulnerability exploitation in the wild, with the absolute risk reduction (ARR) attributable to PSIRT presence of about 17%.
- *Practical significance of the findings*: While a 17% ARR may appear modest in isolation, it becomes highly significant when contextualised within environments characterised by large attack surfaces, widespread software deployment, or critical

infrastructure use. For major software vendors, cloud providers, and industrial suppliers, this reduction translates into a meaningful mitigation of operational, financial, reputational, and societal risks.

- *Impact of confounding factors*: The availability of POC is confirmed as a major external risk factor, altering the ARR by approximately 10%. Industry sector and open-source product characteristics have smaller but still relevant effects. Importantly, these results validate the model's ability to separate organisational practices from external drivers of exploitation risk.
- *Impact on offensive cyber operations*: By reducing the likelihood and predictability of successful exploitation, vulnerability management driven by PSIRTs increases uncertainty for attackers and raises the operational costs of OCO. In particular for state-backed and advanced persistent threat (APT) actors, failed exploitation attempts imply lost vulnerabilities, wasted preparation effort, and potential political or reputational consequences.

Taken together, the findings support the general hypothesis that cyber armament and politically motivated offensive cyber operations can be constrained – though not eliminated – through the systematic adoption of security-by-design and vulnerability management practices by suppliers of digital products. PSIRT, as an organisational mechanism, plays a central role in operationalising these practices and linking technical remediation to coordinated disclosure and ecosystem-level information sharing.

At the same time, the analysis highlights important structural limits. The uneven global distribution of PSIRTs – concentrated among large suppliers in technologically advanced or strategically invested states – creates persistent weak points in the global supply chain. Smaller suppliers, organisations in developing economies, and much of the open-source ecosystem remain comparatively exposed. Moreover, patching constraints in critical infrastructure environments, fragmented regulatory regimes, and geopolitical tensions limit the cumulative effect of good practices.

Based on the findings, three sets of recommendations can be formulated:

1. *Suppliers of digital* products should treat PSIRT as a core organisational capability rather than an optional add-on, and couple it with other security-by-design and secure development lifecycle (SDL) practices, including supply-chain governance, SBOM adoption, and vulnerability management. They should also use empirical methodologies such as the one proposed in this dissertation to inform evidence-based decisions about practices.
2. *Policymakers and regulators* should require vulnerability management capabilities at the supplier level, develop a trusted coordinated vulnerability disclosure (CVD) policy environment, and promote regulatory harmonisation and international cooperation to avoid fragmentation of supply-chain security obligations.
3. *Academic researchers* should apply and extend the proposed assessment model to other security-by-design practices, including secure coding, memory-safe languages, bug-bounty programmes, SBOM governance, and AI-assisted security tools. They should

also Integrate richer datasets on exploited vulnerabilities, APT activity, and supply-chain relationships as they become available. This will enable them to explore combinations of practices and their cumulative effects.

The dissertation makes several distinct contributions to the academic field:

- It introduces and empirically tests a novel, reproducible model for assessing the effectiveness of organisational security practices in reducing real-world vulnerability exploitation.
- It provides the first systematic empirical evaluation of PSIRT effectiveness, using open datasets and statistically grounded methods.
- It maps and analyses the global distribution of more mature PSIRTs, revealing structural inequalities in security capabilities across regions and industries.
- It bridges technical cybersecurity research with international security studies by linking vulnerability management practices to the feasibility and cost of offensive cyber operations.
- It advances interdisciplinary thinking by connecting technological, organisational, economic, and geopolitical dimensions of digital security.

The study has several limitations that also point to productive avenues for future research. The availability and quality of public data on details of the deployed industry practices, including PSIRTs, remain limited, and greater standardisation and transparency would enable more granular analysis of maturity levels and internal structures. Larger and more stratified datasets could strengthen statistical confidence and allow sector-specific insights relevant for studying the impact on OCO.

Future research could also extend the model to assess:

- cost–benefit trade-offs of PSIRT implementation
- comparative efficiency of different vulnerability management practices
- effects of automation and AI-supported tooling
- longer-term strategic impacts on vulnerability markets, stockpiling behaviour, and international peace and security

In conclusion, this dissertation demonstrates that supplier-level security practices – and PSIRT in particular – have strategic implications for the global cyberthreat and international security landscape. While good practices alone cannot halt cyber armament, they can measurably increase the cost, complexity, and uncertainty of offensive cyber operations. In doing so, they represent a technically grounded and politically relevant lever for strengthening digital security in an increasingly contested cyberspace.

Literature

1. Ablon, L., & Bogart, A. (2017). *Zero days, thousands of nights*. Santa Monica, CA: RAND Corporation. Available at: https://www.rand.org/pubs/research_reports/RR1751.html (accessed 25 November 2025).
2. Admass, W.S., Munaye, Y.Y., & Diro, A.A. (2024). Cyber security: State of the art, challenges and future direction. *Cyber Security and Applications*, 2, 100031. doi:10.1016/j.csa.2023.100031
3. Adobe. (2018). *Secure engineering white paper*. Adobe. Available at: https://www.adobe.com/content/dam/cc/en/security/pdfs/adobe-secure-engineering-wp.pdf (accessed 25 November 2025).
4. Ahmad, A., Webb, J., Desouza, K.C., & Boorman, J. (2019). Strategically-motivated advanced persistent threat: Definition, process, tactics and a disinformation model of counterattack. *Computers & Security*, 86, 402–18. doi:10.1016/j.cose.2019.07.001
5. Airike, P-E., Rotter, J.P., & Mark-Herbert, C. (2016). Corporate motives for multi-stakeholder collaboration – corporate social responsibility in the electronics supply chains. *Journal of Cleaner Production*, 131, 639–48. doi:10.1016/j.jclepro.2016.04.121
6. Akinsuli, O. (2021). The rise of AI-enhanced ransomware-as-a-service (RaaS): A new threat frontier. *World Journal of Advanced Engineering Technology and Sciences*, 1(2), 85–97. doi:10.30574/wjaets.2021.1.2.0019
7. Alkhadra, R., Abuzaid, J., AlShammari, M., & Mohammad, N. (2021). SolarWinds hack: In-depth analysis and countermeasures. *Proceedings of the 12th International Conference on Computing Communication and Networking Technologies (ICCCNT)*, 1–7. Kharagpur, India: IEEE. doi:10.1109/ICCCNT51525.2021.9579611
8. Allam, H. (2022). Security-driven pipelines: Embedding DevSecOps into CI/CD workflows. *International Journal of Emerging Trends in Computer Science and Information Technology*, 3(1), 86–97. doi:10.63282/3050-9246.IJETCSIT-V3I1P110
9. Allodi, L., Cremonini, M., Massacci, F., & Shim, W. (2018). The effect of security education and expertise on security assessments: The case of software vulnerabilities. Paper presented at: *17th Annual Workshop on the Economics of Information Security (WEIS)*. 18–19 June 2018; Innsbruck, Austria. doi:10.48550/arXiv.1808.06547
10. Allodi, L., & Massacci, F. (2014). Comparing vulnerability severity and exploits using case-control studies. *ACM Transactions on Information and System Security*, 17(1), 1–20. doi:10.1145/2630069
11. Almukaynizi, M., Nunes, E., Dharaiya, K., Senguttuvan, M., Shakarian, J., et al. (2017). Proactive identification of exploits in the wild through vulnerability mentions online. *Proceedings of the International Conference on Cyber Conflict (CyCon U.S.)*. 7–8 November 2017; Washington, DC. doi:10.1109/CYCONUS.2017.8167501
12. Ali, M., Khan, S., & Vasilakos, A. (2015). Security in cloud computing: Opportunities and challenges. *Information Sciences*, 305, [page range missing]. doi:10.1016/j.ins.2015.01.025

13. Anderson, R. (2001). Why information security is hard – an economic perspective. *Proceedings of the Seventeenth Annual Computer Security Applications Conference*, 358–65. Washington, DC: IEEE Computer Society. doi:10.1109/ACSAC.2001.991552
14. Arbaugh, W.A., Fithen, W.L., & McHugh, J. (2000). Windows of vulnerability: A case study analysis. *Computer*, 33(12), 52–59. doi:10.1109/2.889093
15. Arora, A., Krishnan, R., Nandkumar, A., Telang, R., & Yang, Y. (2004). Impact of vulnerability disclosure and patch availability – an empirical analysis. *Proceedings of the Third Annual Workshop on the Economics of Information Security (WEIS)*, [page range missing].
16. Arora, A., Krishnan, R., Telang, R., & Yang, Y. (2006). An empirical analysis of software vendors' patching behavior: Impact of vulnerability disclosure. In: *International Conference on Information Systems (ICIS)*. Association for Information Systems, [page range missing].
17. Badampudi, D., Wohlin, C., & Petersen, K. (2016). Software component decision-making: In-house, OSS, COTS or outsourcing – a systematic literature review. *Journal of Systems and Software*, 121, 105–24. doi:10.1016/j.jss.2016.07.027
18. Baines, J. (2023). The VulnCheck 2022 exploited vulnerability report – a year-long review of the CISA KEV catalog. *VulnCheck Blog*. Available at: https://vulncheck.com/blog/2022-cisa-kev-review (accessed 25 November 2025).
19. Balsam, A., Nowak, M., Walkowski, M., Oko, J., & Sujecki, S. (2023). Analysis of CVSS vulnerability base scores in the context of exploits' availability. *Proceedings of the 23rd International Conference on Transparent Optical Networks (ICTON)*. Bucharest, Romania, 2–6 July 2023. New York, NY: IEEE, 1–4. doi:10.1109/ICTON59386.2023.10207394
20. Baram, G. (2023). Deter, disrupt, or deceive: Assessing cyber conflict as an intelligence contest. *International Studies Review*, 25(4), viad047. doi:10.1093/isr/viad047
21. Benjamin, V., Li, W., Holt, T., & Chen, H. (2015). Exploring threats and vulnerabilities in hacker web: Forums, IRC and carding shops. *Proceedings of the IEEE International Conference on Intelligence and Security Informatics*, 85–90. doi:10.1109/ISI.2015.7165944
22. Black Duck. (2025). *BSIMM15 report*. April 2025. Available at: https://www.blackduck.com/content/dam/black-duck/en-us/reports/bsimm-report.pdf (accessed 25 November 2025).
23. Böhme, R. (2006). A comparison of market approaches to software vulnerability disclosure. In: Müller, G. (ed.), *Emerging Trends in Information and Communication Security*. Lecture Notes in Computer Science, vol. 3995. Berlin: Springer, 298–311. doi:10.1007/11766155_21
24. Brundage, M., et al. (2018). The malicious use of artificial intelligence: Forecasting, prevention, and mitigation. *arXiv preprint*. Available at: https://arxiv.org/abs/1802.07228 (accessed 25 November 2025).
25. Brunner, E.M., & Suter, M. (2008). *International CIIP handbook 2008/2009*. Zurich: Center for Security Studies, ETH Zurich. Available at:

http://www.css.ethz.ch/publications/pdfs/CIIP-HB-08-09.pdf (accessed 25 November 2025).

26. Buchheit, M., Hermeling, M., Hirsch, F., Martin, B., & Rix, S. (2020). *Software trustworthiness best practices*. Industrial Internet Consortium white paper. Available at: https://www.iiconsortium.org/pdf/Software_Trustworthiness_Best_Practices_Whitepaper_2020_03_23.pdf (accessed 25 November 2025).
27. Bulut, M.F., Adebayo, A., Sow, D., & Ocepek, S. (2022). Vulnerability prioritization: An offensive security approach. *Computing Research Repository (CoRR)*, arXiv:2206.11182. doi:10.48550/arXiv.2206.11182
28. Bunke, M., & Sohr, K. (2011). An architecture-centric approach to detecting security patterns in software. In: Erlingsson, Ú., Wieringa, R., & Zannone, N. (eds.), *Engineering Secure Software and Systems*. Lecture Notes in Computer Science, vol. 6542. Berlin: Springer. doi:10.1007/978-3-642-19125-1_12
29. Charrier, C., & Semrau, J. (2023). Analysis of time-to-exploit trends: 2021–2022. *Mandiant*. Available at: https://cloud.google.com/blog/topics/threat-intelligence/time-to-exploit-trends-2021-2022 (accessed 25 November 2025).
30. Charrier, C., & Weiner, R. (2024). How low can you go? An analysis of 2023 time-to-exploit trends. *Google Threat Intelligence*, 15 October. Available at: https://cloud.google.com/blog/topics/threat-intelligence/time-to-exploit-trends-2023 (accessed 25 November 2025).
31. Charter of Trust. (2018). *Charter of trust: Our 10 principles*. Available at: https://www.charteroftrust.com/about/ (accessed 25 November 2025).
32. Charter of Trust. (2019). *Achieving security by default for products, functionalities, and technologies – baseline requirements*. Available at: https://www.charteroftrust.com/wp-content/uploads/2020/05/200212-P3-Phase-1-Baseline-Requirements_FINAL.pdf (accessed 25 November 2025).
33. Charter of Trust. (2020). *Common risk-based approach for the digital supply chain*. Available at: https://www.charteroftrust.com/wp-content/uploads/2020/02/20-02-11-CoT-P2-phase-1-report.pdf (accessed 25 November 2025).
34. Chen, P., Desmet, L., & Huygens, C. (2014). A study on advanced persistent threats. In: Decker, B.D., & Zúquete, A. (eds.), *Communications and Multimedia Security*. Lecture Notes in Computer Science, vol. 8735. Springer, 63–72. doi:10.1007/978-3-662-44885-4_5
35. Christey, S., & Wysopal, S. (2002). Responsible vulnerability disclosure process. *Internet-Draft*. Internet Engineering Task Force, 15 February. Available at: https://datatracker.ietf.org/doc/draft-christey-wysopal-vuln-disclosure/ (accessed 25 November 2025).
36. Cisco. (2016). *Secure development lifecycle*. Cisco. Available at: https://www.cisco.com/c/dam/en_us/about/doing_business/trust-center/docs/cisco-secure-development-lifecycle.pdf (accessed 25 November 2025).
37. Computer Fraud & Security. (2020). Ransomware claims first fatality as healthcare under renewed assault. *Computer Fraud & Security*, 2020(10), 1–3. doi:10.1016/S1361-3723(20)30101-9

38. Coward, M. (2009). Network-centric violence, critical infrastructure and the urbanization of security. *Security Dialogue*, 40(4–5), 399–418.
39. Crowley, C., & Pescatore, J. (2019). *Common and best practices for security operations centers: Results of the 2019 SOC survey*. Bethesda, MD: SANS. Available at: https://www.sans.org/media/analyst-program/common-practices-security-operations-centers-results-2019-soc-survey-39060.pdf (accessed 25 November 2025).
40. Cybellum. (2023). How to build a strong PSIRT: People, processes & tech. Available at: https://cybellum.com/blog/how-to-build-a-strong-psirt-people-processes-tech/ (accessed 25 November 2025).
41. Cyber Security Agency of Singapore (CSAS). (2017). *Security-by-design framework*. Singapore: Cyber Security Agency of Singapore, 9 November. Available at: https://www.csa.gov.sg/legislation/supplementary-references (accessed 25 November 2025).
42. Cyber Security Agency of Singapore (CSAS). (2025). About cybersecurity labelling scheme for IoT – CLS(IoT). Available at: https://www.csa.gov.sg/our-programmes/certification-and-labelling-schemes/cybersecurity-labelling-scheme/about/ (accessed 25 November 2025).
43. Cybersecurity and Infrastructure Security Agency (CISA). (2021). *Binding operational directive 22-01: Reducing the significant risk of known exploited vulnerabilities*. United States Department of Homeland Security, 3 November.
44. Cybersecurity and Infrastructure Security Agency (CISA). (2022). 2021 trends show increased globalized threat of ransomware. *CISA Advisories*, 10 February. Available at: https://www.cisa.gov/news-events/cybersecurity-advisories/aa22-040a (accessed 25 November 2025).
45. Cybersecurity Tech Accord. (2018). *The cybersecurity tech accord*. Available at: https://cybertechaccord.org/accord/ (accessed 25 November 2025).
46. Cyberspace Solarium Commission. (2020). *Final report*. Washington, DC, March 2020. Available at: https://www.solarium.gov/report (accessed 25 November 2025).
47. Devaux, T., Massip, T., Ulliac, A., Simoni, J.-L., & Varela, P. (2021). Automation of risk-based vulnerability management based on a cyber kill chain model. Paper presented at: *Computer Electronics Security Application Rendezvous (C&ESAR'21)*. 16–17 November 2021; Rennes, France.
48. Dietrich, J., Rasheed, S., Jordan, A., & White, T. (2024). On the security blind spots of software composition analysis. *Proceedings of the ACM Conference* [conference name missing], [page range missing]. doi:10.1145/3689944.3696165
49. Dodson, D., Souppaya, M., & Scarfone, K. (2020). *Mitigating the risk of software vulnerabilities by adopting a secure software development framework (SSDF)*. Gaithersburg, MD: National Institute of Standards and Technology. doi:10.6028/NIST.CSWP.04232020
50. Doll, R., & Hill, A.B. (1950). Smoking and carcinoma of the lung. *British Medical Journal*, 2(4682), 739–48.
51. Dong, J., Chen, S., Ding, F., et al. (2025). Spatiotemporal characteristics and drivers of global cyber conflicts. *Humanities and Social Sciences Communications*, 12, 665. doi:10.1057/s41599-025-04897-7

52. Dougherty, C.R., Sayre, K., Seacord, R.C., Svoboda, D., & Togashi, K. (2009). *Secure design patterns*. Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University. doi:10.1184/R1/6583640.V1
53. Dunn Cavelty, M. (2014). Breaking the cyber-security dilemma: Aligning security needs and removing vulnerabilities. *Science and Engineering Ethics*, 20, 701–15. doi:10.1007/s11948-014-9551-y
54. Eggenschwiler, J. (2018). *Geneva dialogue on responsible behaviour in cyberspace: Private sector framework document*. Zurich: Geneva Dialogue on Responsible Behaviour in Cyberspace (ETH Zurich). Available at: https://genevadialogue.ch/wp-content/uploads/Geneva-Dialogue-Role-of-the-Private-Sector.pdf (accessed 25 November 2025).
55. Egloff, F.J., & Smeets, M. (2021). Publicly attributing cyber attacks: A framework. *Journal of Strategic Studies*, [volume/issue missing], 1–32. doi:10.1080/01402390.2021.1895117
56. European Union Agency for Cybersecurity (ENISA). (2019). *Good practices for security of IoT: Secure software development lifecycle*. ENISA, November. Available at: https://www.enisa.europa.eu/publications/good-practices-for-security-of-iot-1/at_download/fullReport (accessed 25 November 2025).
57. European Union Agency for Cybersecurity (ENISA). (2023). *Good practices for supply chain cybersecurity*. ENISA, June. Available at: https://www.enisa.europa.eu/publications/good-practices-for-security-of-digital-products (accessed 25 November 2025).
58. Eriksson, J., & Giacomello, G. (2014). International relations, cybersecurity, and content analysis: A constructivist approach. In: *The global politics of science and technology – Vol. 2: Perspectives, cases and methods*. Berlin: Springer, 205–19.
59. European Union (EU). (2020). *EU cybersecurity strategy for the digital decade*. Available at: https://digital-strategy.ec.europa.eu/en/library/eus-cybersecurity-strategy-digital-decade-0 (accessed 25 November 2025).
60. European Union (EU). (2022). Directive (EU) 2022/2555 on measures for a high common level of security of network and information systems across the Union (NIS2 Directive). *Official Journal of the European Union*. Available at: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:02022L2555-20221227 (accessed 25 November 2025).
61. European Union (EU). (2024). Regulation (EU) 2024/284 on horizontal cybersecurity requirements for products with digital elements (Cyber Resilience Act). *Official Journal of the European Union*.
62. Europol. (2023). *ChatGPT: The impact of large language models on law enforcement*. Europol, March 2023. Available at: https://www.europol.europa.eu/cms/sites/default/files/documents/Tech%20Watch%20Flash%20-%20The%20Impact%20of%20Large%20Language%20Models%20on%20Law%20Enforcement.pdf (accessed 25 November 2025).
63. Evans, L. (1986). The effectiveness of safety belts in preventing fatalities. *Accident Analysis & Prevention*, 18(3), 229–41. doi:10.1016/0001-4575(86)90007-2

64. Fagan, M., Megas, K.N., Scarfone, K., & Smith, M. (2020a). *Foundational cybersecurity activities for IoT device manufacturers (NISTIR 8259)*. Gaithersburg, MD: National Institute of Standards and Technology. doi:10.6028/NIST.IR.8259
65. Fagan, M., Megas, K.N., Scarfone, K., & Smith, M. (2020b). *IoT device cybersecurity capability core baseline (NISTIR 8259A)*. Gaithersburg, MD: National Institute of Standards and Technology. doi:10.6028/NIST.IR.8259A
66. Finnemore, M., & Hollis, D.B. (2016). Constructing norms for global cybersecurity. *American Journal of International Law*, 110(3), 425–79.
67. FIRST. (2018). *Common vulnerability scoring system version 3.1: Specification document*. Available at: https://www.first.org/cvss/v3.1/specification-document (accessed 25 November 2025).
68. FIRST. (2019). *Product security incident response team (PSIRT) maturity document*. Available at: https://www.first.org/standards/frameworks/psirts/psirt_maturity_document (accessed 25 November 2025).
69. FIRST. (2020a). *Guidelines and practices for multi-party vulnerability coordination and disclosure*. Available at: https://www.first.org/global/sigs/vulnerability-coordination/multiparty/guidelines-v1.1 (accessed 25 November 2025).
70. FIRST. (2020b). *Product security incident response team (PSIRT) services framework* (version 1.1). Available at: https://www.first.org/standards/frameworks/psirts/FIRST_PSIRT_Services_Framework_v1.1.pdf (accessed 25 November 2025).
71. FIRST. (2024). *Common vulnerability scoring system version 4.0: Specification document* (version 1.2). Available at: https://www.first.org/cvss/v4-0/cvss-v40-specification.pdf (accessed 25 November 2025).
72. Forsberg, J., & Frantti, T. (2023). Technical performance metrics of a security operations center. *Computers & Security*, 135, 103529. doi:10.1016/j.cose.2023.103529
73. Frei, S., Tellenbach, B., & Plattner, B. (2008). 0-day patch exposing vendors' (in)security performance. *Black Hat Europe*. Available at: https://www.blackhat.com/presentations/bh-europe-08/Frei/Whitepaper/bh-eu-08-frei-WP.pdf (accessed 25 November 2025).
74. Fucci, D., Alégroth, E., Felderer, M., & Johannesson, C. (2024). Evaluating software security maturity using OWASP SAMM: Different approaches and stakeholders' perceptions. *Journal of Systems and Software*, [volume missing], 112062. doi:10.1016/S0164121224001079
75. Garrity, P. (2024). VulnCheck exploited vulnerabilities report – May 2024. *VulnCheck Blog*, 13 June. Available at: https://vulncheck.com/blog/kev-report-may-2024 (accessed 25 November 2025).
76. Gartzke, E. (2013). The myth of cyberwar: Bringing war in cyberspace back down to earth. *International Security*, 38(2), 41–73. doi:10.1162/ISEC_a_00136
77. Gavrilović, A., Grottola, S., Ittelson, P., Kazakova, A., Petit Siemens, S., Radunović, V., et al. (2025). *OEWG and cybersecurity negotiations at the United Nations: Issues, dynamics, and divides*. Geneva: DiploFoundation, July 2025. Available at:

https://www.diplomacy.edu/resource/oewg-and-cybersecurity-negotiations-at-the-united-nations-issues-dynamics-and-divides/ (accessed 25 November 2025).
78. Gaurav, M. (2024). Embedding security into the pipeline: A framework for scalable DevSecOps implementation. *International Journal of Data Science and Machine Learning*, 1, 37–62.
79. Ghazouani, M., Medromi, H., Sayouti, A., & Faris, S. (2014). Information security risk assessment: A practical approach with a mathematical formulation of risk. *International Journal of Computer Applications*, [volume/issue missing], [page range missing]. doi:10.5120/18097-9155
80. Goddy, U., Osimen, G., Fulani, O., Chidozie, F., & Dada, D. (2024). The weaponisation of artificial intelligence in modern warfare: Implications for global peace and security. *Research Journal in Advanced Humanities*, 5, 24–36. doi:10.58256/g2p9tf63
81. Godwin, J.B. III, Kulpin, A., Rauscher, K.F., & Yaschenko, V. (2014). *Russia–U.S. bilateral on cybersecurity: Critical terminology foundations 2*. New York: EastWest Institute and Information Security Institute of Moscow State University. Available at: https://www.eastwest.ngo/idea/critical-terminology-foundations-2 (accessed 25 November 2025).
82. González-Manzano, L., Fuentes, J.M., Ramos, C., Sánchez, Á., & Quispe, F. (2022). Identifying key relationships between nation-state cyberattacks and geopolitical and economic factors: A model. *Security and Communication Networks*, 2022, 5784674.
83. Google. (2023). *Year in review of zero-days*. Google Project Zero. Available at: https://googleprojectzero.blogspot.com/2023/ (accessed 25 November 2025).
84. Guembe, B., Azeta, A., Misra, S., Osamor, V.C., Fernandez-Sanz, L., & Pospelova, V. (2022). The emerging threat of AI-driven cyber attacks: A review. *Applied Artificial Intelligence*, 36, [page range missing]. doi:10.1080/08839514.2022.2037254
85. Gunduz, M.Z., & Das, R. (2020). Cyber-security on smart grid: Threats and potential solutions. *Computer Networks*, 169, 107094. doi:10.1016/j.comnet.2019.107094
86. Hamer, S., Bowen, J., Haque, M.N., Hines, R., Madden, C., & Williams, L. (2025). Closing the chain: How to reduce your risk of being SolarWinds, Log4j, or XZ Utils. *arXiv preprint*. Available at: https://arxiv.org/abs/2503.12192 (accessed 25 November 2025).
87. Hammi, B., & Zeadally, S. (2023). Software supply-chain security: Issues and countermeasures. *Computer*, 56(7), 54–66. doi:10.1109/MC.2023.3273491
88. Hathaway, M.E., & Savage, J.E. (2012). *Stewardship of cyberspace: Duties for Internet service providers*. Toronto: Canada Centre for Global Security Studies, Munk School of Global Affairs, University of Toronto. Available at: https://www.belfercenter.org/sites/default/files/files/publication/cyberdialogue2012_hathaway-savage.pdf (accessed 25 November 2025).
89. Hathaway, M. (2019). *Patching our digital future is unsustainable and dangerous*. CIGI Paper No. 219. Available at: https://www.cigionline.org/sites/default/files/documents/Paper%20no.219web.pdf (accessed 25 November 2025).

90. Henriksen, A. (2019). The end of the road for the UN GGE process: The future regulation of cyberspace. *Journal of Cybersecurity*, 5(1). doi:10.1093/cybsec/tyy009
91. Herpig, S., & Paulus, A. (2024). The Pall Mall process on cyber intrusion capabilities. *Lawfare*, 19 March. Available at: https://www.lawfaremedia.org/article/the-pall-mall-process-on-cyber-intrusion-capabilities (accessed 25 November 2025).
92. Herr, T. (2016). Malware counter-proliferation and the Wassenaar Arrangement. *Proceedings of the 8th International Conference on Cyber Conflict*, 175–90. Tallinn: IEEE.
93. Herr, T. (2017). *Countering the proliferation of malware: Targeting the vulnerability lifecycle*. Cambridge, MA: Belfer Center, Harvard Kennedy School. Available at: https://www.belfercenter.org/sites/default/files/files/publication/CounteringProliferationofMalware.pdf (accessed 25 November 2025).
94. Householder, A.D., Wassermann, G., Manion, A., & King, C. (2017). *The CERT guide to coordinated vulnerability disclosure*. Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University. Available at: https://resources.sei.cmu.edu/library/asset-view.cfm?assetid=503330 (accessed 25 November 2025).
95. Householder, A., Wassermann, G., Manion, A., & King, C. (2020). *CERT® guide to coordinated vulnerability disclosure*. Carnegie Mellon University. doi:10.1184/R1/12367340.v1
96. Howard, M. (2005). A look inside the security development lifecycle at Microsoft. *MSDN Magazine*, November. Available at: https://docs.microsoft.com/en-us/archive/msdn-magazine/2005/november/a-look-inside-the-security-development-lifecycle-at-microsoft (accessed 25 November 2025).
97. Howard, M., & Lipner, S. (2006). *The security development lifecycle*. Redmond, WA: Microsoft Press. doi:10.1007/s11623-010-0021-7
98. Hunter, L.Y., Albert, C.D., Garrett, E., & Rutland, J. (2022). Democracy and cyberconflict: How regime type affects state-sponsored cyberattacks. *Journal of Cyber Policy*, 7, 72–94.
99. Hutchins, E.M., Cloppert, M.J., & Amin, R.M. (2011). Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains. *Leading Issues in Information Warfare & Security Research*, 1, 80. Available at: https://www.lockheedmartin.com/content/dam/lockheed-martin/rms/documents/cyber/LM-White-Paper-Intel-Driven-Defense.pdf (accessed 25 November 2025).
100. IBM. (2025). *IBM X-Force 2025 threat intelligence index*. Available at: https://www.ibm.com/thought-leadership/institute-business-value/en-us/report/2025-threat-intelligence-index (accessed 25 November 2025).
101. International Committee of the Red Cross (ICRC). (2020). Digital technologies and war. *International Review of the Red Cross*, 102(913), 481–492. doi:10.1017/S1816383120000478
102. International Organization for Standardization / International Electrotechnical Commission (ISO/IEC). (2018). *Information technology – Security techniques –*

Information security management systems – Overview and vocabulary (ISO/IEC 27000:2018). Geneva: ISO.

103. International Organization for Standardization (ISO). (2014). *ISO/IEC 29147:2014 – Vulnerability disclosure*. Geneva: ISO. Available at: https://www.iso.org/standard/45170.html (accessed 25 November 2025).
104. International Organization for Standardization / International Electrotechnical Commission (ISO/IEC). (2019). *Information technology – Vulnerability handling processes (ISO/IEC 30111:2019)*. Geneva: International Organization for Standardization.
105. Ittelson, P., Radunović, V., Kazakova, A., Petit Siemens, S., & Gavrilović, A. (2025). UN OEWG concludes, paving way for permanent cybersecurity mechanism. *Digital Watch*, 17 July. Available at: https://dig.watch/updates/un-oewg-concludes-paving-way-for-permanent-cybersecurity-mechanism (accessed 25 November 2025).
106. Jacobs, J., Romanosky, S., Suciu, O., Edwards, B., & Sarabi, A. (2023). Enhancing vulnerability prioritization: Data-driven exploit predictions with community-driven insights. *Proceedings of the IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, 194–206. doi:10.1109/EuroSPW59978.2023.00027
107. Jada, I., & Mayayise, T.O. (2024). The impact of artificial intelligence on organisational cyber security: An outcome of a systematic literature review. *Data and Information Management*, 8(2), 100063. doi:10.1016/j.dim.2023.100063
108. Juniper. (2015). *Juniper Networks secure development lifecycle: Six practices for improving product security*. Juniper Networks. Available at: https://www.juniper.net/assets/us/en/local/pdf/whitepapers/2000526-en.pdf (accessed 25 November 2025).
109. Jung, B., Li, Y., & Bechor, T. (2022). CAVP: A context-aware vulnerability prioritization model. *Computers & Security*, 116, 102639. doi:10.1016/j.cose.2022.102639
110. Kamara, I., & Boom, J. v.d. (2022). *Computer security incident response teams in the reformed Network and Information Security Directive: Good practices*. Tilburg: Tilburg Institute for Law, Technology, and Society (TILT), Tilburg Law School. Available at: https://www.ncsc.nl/binaries/ncsc/documenten/publicaties/2022/oktober/13/index/NCSC_NIS2_D1_Final.pdf (accessed 25 November 2025).
111. Kasuya, M., Morita, N., Yamazaki, H., & Sakiyama, K. (2023). Automated extraction for vulnerability management on PSIRT: A case study of the automotive industry. *Proceedings of the IEEE International Symposium on Computing and Networking Workshops (CANDARW)*, 232–38. doi:10.1109/CANDARW60564.2023.00046
112. Kaufmann, C. (2016). Multistakeholder participation in cyberspace. *Swiss Review of International and European Law*, 26(2), 217–34.
113. Kazakova, A., Radunović, V., & Droz, S. (2023). *Geneva manual on responsible behaviour in cyberspace: Implementation of norms of responsible behaviour in cyberspace by relevant non-state stakeholders (Chapter 1)*. Geneva: Geneva Dialogue on Responsible Behaviour in Cyberspace, DiploFoundation. Available at: https://genevadialogue.ch/geneva-manual/chapter-1/ (accessed 25 November 2025).

114. Kazakova, A., Radunović, V., & Droz, S. (2025). *Geneva manual on responsible behaviour in cyberspace: Implementation of agreed norms and confidence-building measures by relevant non-state stakeholders to protect critical infrastructure (Chapter 2)*. Geneva: Geneva Dialogue on Responsible Behaviour in Cyberspace, DiploFoundation. Available at: https://genevadialogue.ch/geneva-manual/chapter-2/ (accessed 25 November 2025).
115. Khan, R.A., Khan, S.U., Ilyas, M., & Idris, M.Y. (2020). The state of the art on secure software engineering: A systematic mapping study. *Proceedings of the Evaluation and Assessment in Software Engineering (EASE '20)*. New York, NY: ACM, 487–92. doi:10.1145/3383219.3383290
116. Kol, M., & Oberman, S. (2020). *Ripple20: CVE-2020-11896, CVE-2020-11898 – Information leakage*. Technical white paper. JSOF. Available at: https://www.jsof-tech.com/wp-content/uploads/2020/06/JSOF_Ripple20_Technical_Whitepaper_June20.pdf (accessed 25 November 2025).
117. Kudriavtseva, A., & Gadyatskaya, O. (2024). You cannot improve what you do not measure: A triangulation study of software security metrics. *Proceedings of the ACM Symposium on Applied Computing*, 1223–32. doi:10.1145/3605098.3635892
118. Ladd, D., Simorjay, F., Pulikkathara, G., Jones, J., Miller, M., Lipner, S., et al. (2011). *The SDL progress report*. Microsoft. Available at: http://download.microsoft.com/download/c/e/f/cefb7bf3-de0c-4dcb-995a-c1c69659bf49/sdlprogressreport.pdf (accessed 25 November 2025).
119. Lavorgna, A. (2023). Unpacking the political-criminal nexus in state-cybercrimes: A macro-level typology. *Trends in Organized Crime*. doi:10.1007/s12117-023-09486-1
120. Lee, C.J. (2014). An enhanced risk formula for software security vulnerabilities. *ISACA Journal*, July. Available at: https://www.isaca.org/resources/isaca-journal/past-issues/2014/an-enhanced-risk-formula-for-software-security-vulnerabilities (accessed 25 November 2025).
121. Lehto, M. (2022). APT cyberattack modelling: Building a general model. In: Griffin, R.P., Tatarand, U., & Yankson, B. (eds.), *Proceedings of the 17th International Conference on Cyber Warfare and Security*. Reading, UK: Academic Conferences International, 121–29. doi:10.34190/iccws.17.1.36
122. Lemay, A., Calvet, J., Menet, F., & Fernandez, J.M. (2017). Survey of publicly available reports on advanced persistent threat actors. *Computers & Security*, 72, 26–59. doi:10.1016/j.cose.2017.08.005
123. Libicki, M.C. (2007). *Conquest in cyberspace: National security and information warfare*. Cambridge: Cambridge University Press.
124. Libicki, M.C. (2016). *Cyberspace in peace and war*. Annapolis, MD: Naval Institute Press.
125. Lipner, S. (2004). The trustworthy computing security development lifecycle. *Proceedings of the 20th Annual Computer Security Applications Conference*, 1–12. doi:10.1109/CSAC.2004.41
126. Liu, S., Zhou, J., Hu, X., Roseiro Cogo, F., Xia, X., & Yang, X. (2025). An empirical study on vulnerability disclosure management of open source software systems. *ACM*

Transactions on Software Engineering and Methodology, 34(7), Article 214. doi:10.1145/3716822

127. Lund, M.S. (2024). Hybrid threats in cyberspace. In: *Cyberwarfare*. Cham: Palgrave Macmillan. doi:10.4324/9781032617916-7
128. Lupovici, A. (2021). The dog that did not bark, the dog that did bark, and the dog that should have barked: A methodology for cyber deterrence research. *International Studies Review*, 23(4), 1672–98. doi:10.1093/isr/viab032
129. Marandi, M., Bertia, A., & Silas, S. (2023). Implementing and automating security scanning to a DevSecOps CI/CD pipeline. *Proceedings of the World Conference on Communication & Computing (WCONF)*, 1–6. doi:10.1109/WCONF58270.2023.10235015
130. Maschmeyer, L. (2022). A new and better quiet option? Strategies of subversion and cyber conflict. *Journal of Strategic Studies*, 46(3), 570–94. doi:10.1080/01402390.2022.2104253
131. Matwyshyn, A.M. (2010). Hidden engines of destruction: The reasonable expectation of code safety and the duty to warn in digital products. *Florida Law Review*, 62(1), 109–58.
132. Matzkin, S., Shandler, R., & Canetti, D. (2023). The limits of cyberattacks in eroding political trust: A tripartite survey experiment. *British Journal of Politics and International Relations*, 26(4), 1033–54. doi:10.1177/13691481231210383
133. Maxwell, P., & Barnsby, R. (2019). Insecure at any bit rate: Why Ralph Nader is the true OG of the software design industry. *Journal of Cyber Policy*, 4(3), 346–61.
134. McDermott, J., & Fox, C. (1999). Using abuse case models for security requirements analysis. *Proceedings of the 15th Annual Computer Security Applications Conference*, 55–64. doi:10.1109/CSAC.1999.816013
135. McGraw, G. (2006). *Building secure software*. Boston, MA: Addison-Wesley.
136. Mell, P., Bojanova, I., & Galhardo, C.E. (2024). Measuring the exploitation of weaknesses in the wild. *IT Professional*, 26(3), 14–21. doi:10.48550/arXiv.2405.01289
137. Microsoft. (2017). *Digital Geneva Convention proposal*. Available at: https://blogs.microsoft.com/on-the-issues/2017/02/14/need-digital-geneva-convention/ (accessed 25 November 2025).
138. Microsoft. (2025). *Secure future initiative: Progress report*. Microsoft, November 2025. Available at: https://cdn-dynmedia-1.microsoft.com/is/content/microsoftcorp/microsoft/bade/documents/products-and-services/en-us/security/sfi-nov-2025-progress-report.pdf (accessed 25 November 2025).
139. Migues, S., Steven, J., & Ware, J. (2020). *BSIMM 10*. Available at: https://www.bsimm.com/framework.html (accessed 25 November 2025).
140. Miller, C. (2007). The legitimate vulnerability market: The secretive world of 0-day exploit sales. *Proceedings of the Workshop on the Economics of Information Security (WEIS)*. Available at: http://weis2007.econinfosec.org/papers/29.pdf (accessed 25 November 2025).
141. Mitchell, G.E., & Pytlak, A. (2020). Correlates of state-sponsored cyber conflict. In: *Routledge handbook of international cybersecurity*. Abingdon: Routledge, 22–35.

142. MITRE. (2018). *MITRE ATT&CK: Design and philosophy*. McLean, VA: The MITRE Corporation. Available at: https://attack.mitre.org/matrices/enterprise/ (accessed 25 November 2025).
143. Mkhwanazi, T., Shibambu, A., Nefale, V., Mtsweni, J., Phahlamohlaka, J., Mutemwa, M., et al. (2025). Potential cyber threats to the national elections in the digital age in Africa. In: *Proceedings*, [publisher/place missing]. doi:10.1007/978-3-031-81570-6_22
144. Mladenović, D., & Radunović, V. (2018). Defining offensive cyber capabilities. In: *Briefings and memos from the Research Advisory Group*, GCSC Issue Brief 2, Memo 4. The Hague: The Hague Centre for Strategic Studies, 91–134.
145. Morrison, P., Smith, B.H., & Williams, L. (2017). Measuring security practice use: A case study at IBM. *Proceedings of the IEEE/ACM International Conference on Software Engineering Workshops (CESI@ICSE)*, 16–22. doi:10.1109/CESI.2017.4
146. Morgus, R., Smeets, M., & Herr, T. (2017). Countering the proliferation of offensive cyber capabilities. *Global Commission on the Stability of Cyberspace Issue Brief*, 1, 161–87.
147. Mouratidis, H., Giorgini, P., & Manson, G. (2005). When security meets software engineering: A case of modelling secure information systems. *Information Systems*, 30(8), 609–29. doi:10.1016/j.is.2004.06.002
148. Munich Security Conference (MSC). (2015). *Munich Security Report 2015: Collapsing order, reluctant guardians?* Munich: Munich Security Conference. Available at: https://securityconference.org/en/publications/munich-security-report-2015/ (accessed 25 November 2025).
149. Mothanna, Y., Elmedany, W., Hammad, M., Ksantini, R., & Sharif, M.S. (2024). Adopting security practices in software development process: Security testing framework for sustainable smart cities. *Computers & Security*, 144, 103985. doi:10.1016/j.cose.2024.103985
150. Nappa, A., Johnson, R., Bilge, L., Caballero, J., & Dumitras, T. (2015). The attack of the clones: A study of the impact of shared code on vulnerability patching. *Proceedings of the IEEE Symposium on Security and Privacy*. San Jose, CA, 18–20 May 2015, 692–708. doi:10.1109/SP.2015.48
151. North Atlantic Treaty Organization Cooperative Cyber Defence Centre of Excellence (NATO). (2016). *The cyber attack on Ukraine's power grid: A case study*. Available at: https://www.nato.int/cps/en/natolive/official_texts_133169.htm (accessed 25 November 2025).
152. National Initiative for Cybersecurity Careers and Studies (NICCS). (n.d.). *Explore terms: A glossary of common cybersecurity terminology*. U.S. Department of Homeland Security. Available at: http://niccs.us-cert.gov/glossary (accessed 25 November 2025).
153. National Institute of Standards and Technology (NIST). (n.d.). *National Vulnerability Database (NVD) dashboard*. Available at: https://nvd.nist.gov/general/nvd-dashboard (accessed 21 October 2025).
154. National Vulnerability Database News (NVD News). (2024). National Vulnerability Database general update. National Institute of Standards and Technology. Available at: https://www.nist.gov/itl/nvd/nvd-news (accessed 25 November 2025).

155. Organisation for Economic Co-operation and Development (OECD). (2018). *Roles and responsibilities of actors: Governance of digital security in organisations and security of digital technologies*. Paris: OECD. Available at: http://www.oecd.org/officialdocuments/publicdisplaydocumentpdf/?cote=DSTI/CDEP/SPDE(2019)4/FINAL&docLanguage=En (accessed 25 November 2025).
156. Organisation for Economic Co-operation and Development (OECD). (2022). *Recommendation of the Council on the digital security of products and services*. Paris: OECD. Available at: https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0481 (accessed 25 November 2025).
157. Organisation for Economic Co-operation and Development (OECD). (2023). *Good practice guidance on the co-ordination of digital security vulnerabilities*. Paris: OECD, 23 January. Available at: https://one.oecd.org/document/DSTI/CDEP/SDE(2021)9/FINAL/en/pdf (accessed 25 November 2025).
158. O'Donoghue, E., Boles, B., Izurieta, C., & Reinhold, A.M. (2024). Impacts of software bill of materials (SBOM) generation on vulnerability detection. *Proceedings of the ACM Workshop on Software Supply Chain Offensive and Defensive Research (SCORED@CCS)*, 67–76. doi:10.1145/3689944.3696164
159. OWASP. (2019). *SAMM model overview*. OWASP. Available at: https://owaspsamm.org/model/ (accessed 25 November 2025).
160. Palleti, V., Adepu, S., Mishra, V.K., & Mathur, A. (2021). Cascading effects of cyber-attacks on interconnected critical infrastructure. *Cybersecurity*, 4, [article number missing]. doi:10.1186/s42400-021-00071-z
161. Paris Call for Trust and Security in Cyberspace (Paris Call). (2018). *Paris Call: The 9 principles*. 11 December 2018. Available at: https://pariscall.international/en/principles (accessed 25 November 2025).
162. Parsons, D. (2022). *The state of ICS/OT cybersecurity in 2022 and beyond*. Bethesda, MD: SANS, October. Available at: https://www.sans.org/white-papers/state-ics-ot-cybersecurity-2022-beyond/ (accessed 25 November 2025).
163. Patterson, C.M., Nurse, J.R.C., & Franqueira, V.N.L. (2023). Learning from cyber security incidents: A systematic review and future research agenda. *Computers & Security*, 132, 103309. doi:10.1016/j.cose.2023.103309
164. Pawlak P. (2014) *Riding the digital wave: The impact of cyber capacity building on human development*. Institute for Security Studies Paris, December. Available from: https://www.files.ethz.ch/isn/186860/Report_21_Cyber.pdf (accessed 25 November 2025)
165. Perlroth, N. (2021). *This is how they tell me the world ends: The cyberweapons arms race*. New York, NY: Bloomsbury Publishing.
166. Pleskonjić, D., Tica, L., Jelić, V., Todorović, D., & English, A. (2025). AI's role in cybersecurity threats and defenses: Examples of concrete solutions. *Proceedings of INFOTECH 2025*, 75–80. Available at: https://infotech.org.rs/wp-content/uploads/2025/06/INFOTECH-2025-Proceedings-1.pdf (accessed 25 November 2025).

167. Prasad, N., Diro, A., Warren, M., & Fernando, M. (2025). A survey of cyber threat attribution: Challenges, techniques, and future directions. *Computers & Security*, 157, 104606. doi:10.1016/j.cose.2025.104606
168. Przymus, P., & Durieux, T. (2025). Wolves in the repository: A software engineering analysis of the XZ Utils supply chain attack. *Proceedings of the IEEE/ACM International Conference on Mining Software Repositories (MSR)*, 91–102. doi:10.1109/MSR66628.2025.00026
169. Radunović, V. (2017). *Towards a secure cyberspace via regional cooperation*. Geneva: DiploFoundation. Available at: https://www.diplomacy.edu/resource/towards-a-secure-cyberspace-via-regional-co-operation/ (accessed 25 November 2025).
170. Radunović, V., & Gavrilović, A. (2025). Mapping: Offensive cyber capabilities. *Digital Watch Observatory*, DiploFoundation. Available at: https://infogram.com/offensive-cyber-capabilities-1h7z2l87v9njx6o (accessed 26 November 2025).
171. Radunović, V., & Grätz-Hoffmann, J. (2020). *Security of digital products and services: Reducing vulnerabilities and secure design (industry good practices)*. Geneva: Geneva Dialogue on Responsible Behaviour in Cyberspace, DiploFoundation. Available at: https://genevadialogue.ch/goodpractices/ (accessed 25 November 2025).
172. Radunović, V., Grätz-Hoffmann, J., & Maciel, M. (2021). Impact of good corporate practices for security of digital products on global cyber stability. *Proceedings of the 13th International Conference on Cyber Conflict (CyCon)*, 25–42. doi:10.23919/CyCon51939.2021.9467805
173. Radunović, V., Veinović, M., & Jevremović, A. (2025). The efficiency of ICT suppliers' Product Security Incident Response Teams in reducing the risk of exploitation of vulnerabilities in the wild. *Computers & Security*, 152, 104388. doi:10.1016/j.cose.2025.104388
174. Rajnovic, D. (2010). *Computer incident response and product security*. Harlow: Pearson Education.
175. Ramamoorthy, S. (2021). Test automation for application security. *Cisco Blogs*, 5 April. Available at: https://blogs.cisco.com/security/test-automation-for-application-security (accessed 25 November 2025).
176. Ramaswamy, Y. (2023). DevSecOps in practice: Embedding security automation into agile software delivery pipelines. *Journal of Computational Analysis and Applications*, 31, 1372–81. doi:10.48047/jocaaa.2023.31.04.27
177. Ravindrababu, S.G., & Alves-Foss, J. (2023). Analysis of vulnerability trends and attacks in OT systems. In: Yang, X.S., Sherratt, S., Dey, N., & Joshi, A. (eds.), *Proceedings of the Seventh International Congress on Information and Communication Technology*, vol. 2. Singapore: Springer, 127–42. doi:10.1007/978-981-19-1610-6_12
178. Raza, M., Imtiaz, A., & Shoaib, U. (2019). A review on security issues and their impact on hybrid cloud computing environment. *International Journal of Advanced Computer Science and Applications*, 10, 353–57. doi:10.14569/IJACSA.2019.0100346
179. Rid, T., & Buchanan, B. (2014). Attributing cyber attacks. *Journal of Strategic Studies*, 38(1–2), 4–37. doi:10.1080/01402390.2014.977382

180. Rizmal, I., & Radunović, V. (2019). *Baseline study*. Geneva: Geneva Dialogue on Responsible Behaviour in Cyberspace. Available at: https://genevadialogue.ch/wp-content/uploads/Geneva-Dialogue-Baseline-Study.pdf (accessed 25 November 2025).
181. Romanosky, S. (2019). Developing an objective, repeatable scoring system for a vulnerability equity process. *Lawfare*, 4 February. Available at: https://www.lawfareblog.com/developing-objective-repeatable-scoring-system-vulnerability-equities-process (accessed 25 November 2025).
182. Ruohonen, J., & Allodi, L. (2018). A bug bounty perspective on the disclosure of web vulnerabilities. *Proceedings of the Workshop on the Economics of Information Security (WEIS)*. Innsbruck, Austria.
183. Sabottke, C., Suciu, O., & Dumitras, T. (2015). Vulnerability disclosure in the age of social media: Exploiting Twitter for predicting real-world exploits. *Proceedings of the USENIX Security Symposium*, 1041–56.
184. SAFECode. (2017). *Managing security risks inherent in the use of third-party components*. SAFECode. Available at: https://safecode.org/wp-content/uploads/2017/05/SAFECode_TPC_Whitepaper.pdf (accessed 25 November 2025).
185. SAFECode. (2018). *Fundamental practices for secure software development: Essential elements of a secure development lifecycle program* (3rd ed.). SAFECode, March. Available at: https://safecodedev.wpengine.com/wp-content/uploads/2018/03/SAFECode_Fundamental_Practices_for_Secure_Software_Development_March_2018.pdf (accessed 25 November 2025).
186. Santos, J.C.S., Tarrit, K., & Mirakhorli, M. (2017). A catalog of security architecture weaknesses. *Proceedings of the IEEE International Conference on Software Architecture Workshops*, 220–23. doi:10.1109/ICSAW.2017.25
187. Santos, D. dos, Dashevskyi, S., Wetzels, J., & Amri, A. (2020). *Amnesia:33*. Forescout. Available at: https://www.forescout.com/company/resources/amnesia33-how-tcp-ip-stacks-breed-critical-vulnerabilities-in-iot-ot-and-it-devices/ (accessed 25 November 2025).
188. Sauerwein, C., Sillaber, C., Huber, M.M., Mussmann, A., & Breu, R. (2018). The tweet advantage: An empirical analysis of 0-day vulnerability information shared on Twitter. In: Janczewski, L.J., & Kutylowski, M. (eds.), *Security and Privacy Protection in Information Processing Systems*. IFIP Advances in Information and Communication Technology, vol. 529. Cham: Springer, 201–15. doi:10.1007/978-3-319-99828-2_15
189. Shevchenko, N. (2018). Threat modeling: 12 available methods. *SEI Blog*, 3 December. Available at: https://insights.sei.cmu.edu/sei_blog/2018/12/threat-modeling-12-available-methods.html (accessed 25 November 2025).
190. Shourya, R., Kumagai, Y., Yamazaki, H., & Nakakoji, H. (2023). Proposal of vulnerability assessment tool for software supply chain security. *Journal of Information Processing*, 31, 842–50. doi:10.2197/ipsjjip.31.842
191. Shukla, O. (2025). Software supply chain security: Designing a secure solution with SBOM for modern software ecosystems. *International Journal of Engineering Research & Technology*, 14(4).

192. Siemens. (n.d.). *Siemens vulnerability handling and disclosure process*. Siemens. Available at: https://www.siemens.com/global/en/products/services/cert/vulnerability-process.html (accessed 25 November 2025).
193. Sigholm, J., & Larsson, E. (2014). Determining the utility of cyber vulnerability implantation: The Heartbleed bug as a cyber operation. *Proceedings of the IEEE Military Communications Conference*, Baltimore, MD, 110–16. doi:10.1109/MILCOM.2014.25
194. Simon, T., & Haklai, B. (2025). Data disaster – A new threat for sovereignty. *International Journal of Disaster Risk Reduction*, 125, 105538. doi:10.1016/j.ijdrr.2025.105538
195. Sindre, G., & Opdahl, A.L. (2005). Eliciting security requirements with misuse cases. *Requirements Engineering*, 10, 34–44. doi:10.1007/s00766-004-0194-4
196. Sivagnanam, A., Atefi, S., Ayman, A., Grossklags, J., & Laszka, A. (2021). On the benefits of bug bounty programs: A study of Chromium vulnerabilities. *Proceedings of the Workshop on the Economics of Information Security (WEIS)*.
197. Smeets, M. (2018). The strategic promise of offensive cyber operations. *Strategic Studies Quarterly*, 12(3), 90–113.
198. Sophos. (2025). *State of ransomware 2025*. Sophos white paper, June 2025. Available at: https://assets.sophos.com/X24WTUEQ/at/9brgj5n44hqvgsp5f5bqcps/sophos-state-of-ransomware-2024-wp.pdf (accessed 25 November 2025).
199. Spring, J.M. (2023). An analysis of how many undiscovered vulnerabilities remain in information systems. *Computers & Security*, 131, 103191. doi:10.1016/j.cose.2023.103191
200. Squillace, J., Cappella, J., & Sepp, A. (2024). User vulnerabilities in AI-driven systems: Current cybersecurity threat dynamics and malicious exploits in supply chain management and project management. *Proceedings of the ASU International Conference on Emerging Technologies for Sustainability and Intelligent Systems*, 1760–65. doi:10.1109/ICETSIS61505.2024.10459480
201. Stamp, M. (2011). *Information security: Principles and practice*. Hoboken, NJ: John Wiley & Sons.
202. Stedmon, N. (2020). The impact of cybersecurity threats on the 2020 US elections. *arXiv preprint*. doi:10.48550/arXiv.2012.08968
203. Stoddart, K. (2022). Introduction. In: *Cyberwarfare*. Palgrave Studies in Cybercrime and Cybersecurity. Cham: Palgrave Macmillan. doi:10.1007/978-3-030-97299-8_1
204. Stone, M., Semrau, J., & Sadowski, J. (2024). *We're all in this together: A year in review of zero-days exploited in-the-wild in 2023*. Google. Available at: https://storage.googleapis.com/gweb-uniblog-publish-prod/documents/Year_in_Review_of_ZeroDays.pdf (accessed 25 November 2025).
205. Strom, B.E., Applebaum, A., Miller, D.P., Nickels, K.C., Pennington, A.G., & Thomas, C.B. (2018). *MITRE ATT&CK: Design and philosophy*. McLean, VA: The MITRE Corporation.
206. Suby, M., & Dickson, F. (2015). *The 2015 (ISC)² global information security workforce study*. Frost & Sullivan. Available at: https://www.isc2.org/-

/media/Files/Research/GISWS-Archive/GISWS-2015.ashx (accessed 25 November 2025).
207. Syed, R., Rahafrooz, M., & Keisler, J.M. (2018). What it takes to get retweeted: An analysis of software vulnerability messages. *Computers in Human Behavior*, 80, 207–15. doi:10.1016/j.chb.2017.11.024
208. Syed, R. (2020). Cybersecurity vulnerability management: A conceptual ontology and cyber intelligence alert system. *Information & Management*, 57(6), 103334. doi:10.1016/j.im.2020.103334
209. Tahara, J., Watanabe, K., Koshijima, I., & Oka, R. (2025). Resolving conflicts between PSIRT and safety teams: A collaborative approach. *Human Factors in Cybersecurity*, 168, 70. doi:10.54941/ahfe1006144
210. Tanczer, L.M., Brass, I., & Carr, M. (2018). CSIRTs and global cybersecurity: How technical experts support science diplomacy. *Global Policy*, 9(S3), 60–66. doi:10.1111/1758-5899.12625
211. Tomer, S. (2025). Blurring the lines: How nation-states and organized cybercriminals are becoming alike. *Trellix Blog*. Available at: https://www.trellix.com/en-in/blogs/research/blurring-the-lines-how-nation-states-and-cybercriminals-are-becoming-alike/ (accessed 25 November 2025).
212. Trend Micro. (2023). *Everything is connected: Uncovering the ransomware threat from global supply chains*. Trend Micro, January 2023. Available at: https://www.trendmicro.com/explore/glrans (accessed 25 November 2025).
213. United Kingdom Government (UK). (2018a). *Secure development and deployment guidance*. National Cyber Security Centre, 22 November. Available at: https://www.ncsc.gov.uk/collection/developers-collection (accessed 25 November 2025).
214. United Kingdom Government (UK). (2018b). *The equities process*. GCHQ. Available at: https://www.gchq.gov.uk/information/equities-process (accessed 25 November 2025).
215. United Kingdom Government (UK). (2019). Plans announced to introduce new laws for internet connected devices. Press release, 1 May. Available at: https://www.gov.uk/government/news/plans-announced-to-introduce-new-laws-for-internet-connected-devices (accessed 25 November 2025).
216. United Nations (UN). (2013). *Report of the Group of Governmental Experts on Developments in the Field of Information and Telecommunications in the Context of International Security* (A/68/98). Available at: https://undocs.org/A/68/98 (accessed 25 November 2025).
217. United Nations (UN). (2015). *Report of the Group of Governmental Experts on Developments in the Field of Information and Telecommunications in the Context of International Security* (A/70/174). Available at: https://undocs.org/A/70/174 (accessed 25 November 2025).
218. United Nations (UN). (2021). *Group of Governmental Experts on Advancing Responsible State Behaviour in Cyberspace in the Context of International Security* (A/76/135). Available at: https://undocs.org/A/76/135 (accessed 25 November 2025).

219. United Nations (UN). (2025). *Final report of the open-ended working group on security of and in the use of information and communications technologies 2021–2025* (A/AC.292/2025/L.1). Available at: https://undocs.org/A/AC.292/2025/L.1 (accessed 25 November 2025).
220. Uren, T., Hogeveen, B., & Hanson, F. (2018). Defining offensive cyber capabilities. *GCSC Issue Brief 2*. The Hague: The Hague Centre for Strategic Studies, 73–90.
221. Vance, A. (2014). Flow-based analysis of advanced persistent threats detecting targeted attacks in cloud computing. *Proceedings of the International Conference on Problems of Infocommunications Science and Technology*, 173–76. doi:10.1109/INFOCOMMST.2014.6992342
222. Vander Geeten, V. (2024). How can national policies support the development and implementation of coordinated vulnerability disclosure? *Cyber Security: A Peer-Reviewed Journal*, 7(3), 253–61. doi:10.69554/WOER4540
223. Vielberth, M., Böhm, F., Fichtinger, I., & Pernul, G. (2020). Security operations center: A systematic study and open challenges. *IEEE Access*, 8, 227756–79. doi:10.1109/ACCESS.2020.3045514
224. Voicu, L., & Fouque, C. (2025). Time-to-patch metrics: A survival analysis approach using Qualys and Elastic. *Elastic Security Labs*, 22 October. Available at: https://www.elastic.co/security-labs/time-to-patch-metrics (accessed 25 November 2025).
225. Waidner, M., Backes, M., & Müller-Quade, J. (2014). *Development of secure software with security by design*. Darmstadt: Fraunhofer Institute for Secure Information Technology SIT. Available at: https://www.sit.fraunhofer.de/fileadmin/dokumente/studien_und_technical_reports/Trendreport_Security_by_Design_EN.pdf (accessed 25 November 2025).
226. Walshe, T., & Simpson, A. (2020). An empirical study of bug bounty programs. *Proceedings of the IEEE International Workshop on Intelligent Bug Fixing*, 35–44. doi:10.1109/IBF50092.2020.9034828
227. Wee, A., Kudriavtseva, A., & Gadyatskaya, O. (2024). "Have heard of it": A study with practitioners on adoption of secure software development frameworks. *Proceedings of the IEEE European Symposium on Security and Privacy Workshops*, 626–33. doi:10.1109/EuroSPW61312.2024.00076
228. Weir, C., Migues, S., & Williams, L. (2022). Exploring the shift in security responsibility. *IEEE Security & Privacy*, 20(6), 8–17. doi:10.1109/MSEC.2022.3150238
229. Whitehead, D.E., Owens, K., Gammel, D., & Smith, J. (2017). Ukraine cyber-induced power outage: Analysis and practical mitigation strategies. *Proceedings of the Annual Conference for Protective Relay Engineers*, 1–8. doi:10.1109/CPRE.2017.8090056
230. White House. (2017). *Vulnerabilities equities policy and process for the United States Government*. 15 November. Available at: https://trumpwhitehouse.archives.gov/sites/whitehouse.gov/files/images/External%20-%20Unclassified%20VEP%20Charter%20FINAL.PDF (accessed 25 November 2025).

231. White House. (2021). *Executive Order 14028: Improving the Nation's Cybersecurity*. Available at: https://www.nist.gov/itl/executive-order-14028-improving-nations-cybersecurity (accessed 25 November 2025).
232. Williams, L., McGraw, G., & Migues, S. (2018). Engineering security vulnerability prevention, detection, and response. *IEEE Software*, [volume/issue missing], 1–1. doi:10.1109/MS.2018.290110854
233. Wiśniewski, M., Gladysz, B., Ejsmont, K., Wodecki, A., & van Erp, T. (2022). Industry 4.0 solutions impacts on critical infrastructure safety and protection: A systematic literature review. *IEEE Access*, 10, 82716–35. doi:10.1109/ACCESS.2022.3195337
234. World Economic Forum. (2022). *Global cybersecurity outlook 2025*. Available at: https://reports.weforum.org/docs/WEF_Global_Cybersecurity_Outlook_2025.pdf (accessed 25 November 2025).
235. Zhao, M., Grossklags, J., & Liu, P. (2015). An empirical study of web vulnerability discovery ecosystems. *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*, 1105–17. doi:10.1145/2810103.2813704
236. Zhang, S., Zhang, X., & Ou, X. (2014). After we knew it: Empirical study and modeling of cost-effectiveness of exploiting prevalent known vulnerabilities across IaaS cloud. *Proceedings of the ACM Asia Conference on Computer and Communications Security*, 317–28. doi:10.1145/2590296.2590300
237. Zwarts, H., Du Toit, J., & Solms, B. (2025). Augmenting cybersecurity awareness at critical infrastructures in developing countries through a cybersecurity governance maturity model. *Proceedings of the European Conference on Cyber Warfare and Security*, 726–33. doi:10.34190/eccws.24.1.3708

Dataset References

1. [dataset] National Vulnerability Database (NVD) [Internet]. US National Institute of Standards and Technology (NIST). Available from: https://nvd.nist.gov/ (accessed 14 July 2024)
2. [dataset] Known Exploited Vulnerability (KEV) [Internet]. US Cybersecurity and Infrastructure Security Agency (CISA). Available from: https://www.cisa.gov/known-exploited-vulnerabilities-catalog (accessed 14 July 2024)
3. [dataset] Exploit-db (EDB) [Internet]. OffSec. Available from: https://www.exploit-db.com/ (accessed 14 July 2024)

THE IMPACT OF GOOD PRACTICES ON THE SECURITY OF DIGITAL PRODUCTS AND THE LIMITATION OF OFFENSIVE CYBER OPERATIONS

ORIGINALITY REPORT

23%
SIMILARITY INDEX

PRIMARY SOURCES

1	Vladimir Radunović, Mladen Veinović, Aleksandar Jevremović. "The Efficiency of ICT Suppliers' Product Security Incident Response Teams in Reducing the Risk of Exploitation of Vulnerabilities in the Wild", Computers & Security, 2025 Crossref	5687 words — 15%
2	Vladimir Radunović, Mladen Veinović, Aleksandar Jevremović. "The efficiency of ICT suppliers' product security incident response teams in reducing the risk of exploitation of vulnerabilities in the wild", Computers & Security, 2025 Crossref	588 words — 2%
3	www.ccdcoe.org Internet	89 words — < 1%
4	dokumen.pub Internet	70 words — < 1%
5	arxiv.org Internet	60 words — < 1%
6	genevadialogue.ch Internet	57 words — < 1%
	cdn.shopify.com	

7	Internet	50 words — < 1%
8	www.coursehero.com Internet	46 words — < 1%
9	eucd.s3.eu-central-1.amazonaws.com Internet	45 words — < 1%
10	"The Palgrave Handbook on Cyber Diplomacy", Springer Science and Business Media LLC, 2025 Crossref	44 words — < 1%
11	core.ac.uk Internet	36 words — < 1%
12	kth.diva-portal.org Internet	36 words — < 1%
13	www.researchgate.net Internet	36 words — < 1%
14	Klaus-Rainer Müller. "Handbuch Unternehmenssicherheit", Springer Science and Business Media LLC, 2022 Crossref	30 words — < 1%
15	www.cisa.gov Internet	27 words — < 1%
16	ccdcoe.org Internet	26 words — < 1%
17	vulners.com Internet	24 words — < 1%
18	js.vnu.edu.vn Internet	21 words — < 1%

19	www.aseanall.com Internet	21 words — < 1%
20	Lorenzo Bracciale, Sasha Riozzi, Giorgia Panico, Emanuele Raso, Giuseppe Bianchi, Pierpaolo Loreti. "Quantifying medical device cybersecurity risk with CVSS BTE", Scientific Reports, 2025 Crossref	20 words — < 1%
21	eudoxuspress.com Internet	20 words — < 1%
22	ir-library.ku.ac.ke Internet	20 words — < 1%
23	www.interface-eu.org Internet	20 words — < 1%
24	www.mdpi.com Internet	20 words — < 1%
25	Eva Claessen. "Reshaping the internet – the impact of the securitisation of internet infrastructure on approaches to internet governance: the case of Russia and the EU", Journal of Cyber Policy, 2020 Crossref	19 words — < 1%
26	ebin.pub Internet	19 words — < 1%
27	link.springer.com Internet	19 words — < 1%
28	iris.uniroma1.it Internet	18 words — < 1%
29	www.elgaronline.com Internet	18 words — < 1%

30	www.sgpjbg.com Internet	18 words — < 1%
31	"Advances in Production Management Systems. Cyber-Physical-Human Production Systems: Human-AI Collaboration and Beyond", Springer Science and Business Media LLC, 2026 Crossref	16 words — < 1%
32	James Ransome, Brook S. E. Schoenfield. "Building in Security at Agile Speed", CRC Press, 2021 Publications	16 words — < 1%
33	hcss.nl Internet	16 words — < 1%
34	wpkursevi.diplomacy.edu Internet	16 words — < 1%
35	www.nature.com Internet	16 words — < 1%
36	Anthony Elliott. "The Routledge Social Science Handbook of AI", Routledge, 2021 Publications	15 words — < 1%
37	henrykywu-fjtba.wordpress.com Internet	15 words — < 1%
38	www.cisco.com Internet	15 words — < 1%
39	"Security and Trust Management", Springer Science and Business Media LLC, 2026 Crossref	14 words — < 1%
40	library.oapen.org	

Internet

14 words — < 1%

41 www.fastcompany.com
Internet

14 words — < 1%

42 Vladimir Radunovic, Jonas Gratz-Hoffmann, Marilia Maciel. "Impact of Good Corporate Practices for Security of Digital Products on Global Cyber Stability", 2021 13th International Conference on Cyber Conflict (CyCon), 2021
Crossref

13 words — < 1%

43 csirt.fr
Internet

13 words — < 1%

44 www.accesspartnership.com
Internet

13 words — < 1%

45 Eneken Tikk, Mika Kerttunen. "Routledge Handbook of International Cybersecurity", Routledge, 2020
Publications

12 words — < 1%

46 resources.sei.cmu.edu
Internet

12 words — < 1%

47 www.eucybernet.eu
Internet

12 words — < 1%

48 www.securitycouncilreport.org
Internet

12 words — < 1%

49 Emanuele Iannone, Giulia Sellitto, Emanuele Iaccarino, Filomena Ferrucci, Andrea De Lucia, Fabio Palomba. "Early and Realistic Exploitability Prediction of Just-Disclosed Software Vulnerabilities: How Reliable Can It

11 words — < 1%

Be?", ACM Transactions on Software Engineering and Methodology, 2024
Crossref

50	John R. Vacca. "Cloud Computing Security - Foundations and Challenges", CRC Press, 2020 Publications	11 words — < 1%
51	ar5iv.labs.arxiv.org Internet	11 words — < 1%
52	issuu.com Internet	11 words — < 1%
53	mediatum.ub.tum.de Internet	11 words — < 1%
54	standinggroups.ecpr.eu Internet	11 words — < 1%
55	"ISSE 2014 Securing Electronic Business Processes", Springer Science and Business Media LLC, 2014 Crossref	10 words — < 1%
56	Arianna Vettorel. "Rights of Individuals in an Earth Observation and Satellite Navigation Environment", Brill, 2024 Crossref	10 words — < 1%
57	Eider Iturbe, Oscar Llorente-Vazquez, Angel Rego, Erkuden Rios, Nerea Toledo. "Unleashing offensive artificial intelligence: Automated attack technique code generation", Computers & Security, 2024 Crossref	10 words — < 1%
58	Ian Coston, Karl David Hezel, Eadan Plotnizky, Mehrdad Nojoumian. "Enhancing Secure Software Development with AZTRM-D: An AI-Integrated	10 words — < 1%

Approach Combining DevSecOps, Risk Management, and Zero Trust", Applied Sciences, 2025
Crossref

59	Shoki Kosai, Liao Hanqing, Zhengyang Zhang, Kazuyo Matsubae, Eiji Yamasue. "Multi-regional land disturbances induced by mineral use in a product-based approach: A case study of gasoline, hybrid, battery electric and fuel cell vehicle production in Japan", Resources, Conservation and Recycling, 2022 Crossref	10 words — < 1%
60	cyberlaw.ccdcoe.org Internet	10 words — < 1%
61	dash.harvard.edu Internet	10 words — < 1%
62	dig.watch Internet	10 words — < 1%
63	docplayer.net Internet	10 words — < 1%
64	ijmmu.com Internet	10 words — < 1%
65	jyx.jyu.fi Internet	10 words — < 1%
66	lume.ufrgs.br Internet	10 words — < 1%
67	nsarchive.gwu.edu Internet	10 words — < 1%
68	nvlpubs.nist.gov Internet	10 words — < 1%

69	pr.toti.com Internet	10 words — < 1%
70	staging.wisdominterface.com Internet	10 words — < 1%
71	www.oxfordbibliographies.com Internet	10 words — < 1%
72	www.scirp.org Internet	10 words — < 1%
73	www.stiftung-nv.de Internet	10 words — < 1%
74	www.thirdway.org Internet	10 words — < 1%
75	www2.computerworld.com Internet	10 words — < 1%
76	"The Great Power Competition Volume 3", Springer Science and Business Media LLC, 2022 Crossref	9 words — < 1%
77	Chalermpong Senarak. "Port Management Strategy for Enhancing Cybersecurity and Resilience - Learning from Case Studies", Routledge, 2025 Publications	9 words — < 1%
78	Davide Bonaventura, Sergio Esposito, Giampaolo Bella. "A case of smart devices that compromise home cybersecurity", Computers & Security, 2025 Crossref	9 words — < 1%
79	Fergle D'Aubeterre. "An empirical evaluation of information security awareness levels in designing secure business processes", Proceedings of the 4th	9 words — < 1%

International Conference on Design Science Research in Information Systems and Technology - DESRIST 09 DESRIST 09, 2009
Crossref

80 James Ransome, Anmol Misra. "Core Software Security - Security at the Source", Auerbach Publications, 2019 — 9 words — < 1%
Publications

81 Joshua Yindenaba Abor, John Macomber, Thankom Arun, Victor Murinde. "The Routledge Handbook of Infrastructure Finance", Routledge, 2025 — 9 words — < 1%
Publications

82 Luca Allodi, Fabio Massacci. "Comparing Vulnerability Severity and Exploits Using Case-Control Studies", ACM Transactions on Information and System Security, 2014 — 9 words — < 1%
Crossref

83 Maryam Roshanaei, Mark Gregory Jachura. "Integrating AI in Cybersecurity Higher Education: A Path to Workforce Readiness", Journal of Intelligent Learning Systems and Applications, 2025 — 9 words — < 1%
Crossref

84 Michael Enemuo, Olumide Ogunmodimu. "Transitioning the mining sector: A review of renewable energy integration and carbon footprint reduction strategies", Applied Energy, 2025 — 9 words — < 1%
Crossref

85 Mirko Farina, Xiao Yu, Jin Chen. "Digital Development - Technology, Ethics and Governance", Routledge, 2025 — 9 words — < 1%
Publications

86	Muharrem Kılıç, Sezer Bozkus Kahyaoglu. "Digitalization and Women's Rights - Women and Sustainability: Volume 2", CRC Press, 2025 Publications	9 words — < 1%
87	Sai Gautam Mandapati, Chathurika Ranaweera, Robin Doss. "Early Detection and Mitigation of Cache-based Attacks in IoT-NDN", Computers & Security, 2025 Crossref	9 words — < 1%
88	Yunfei Wang, Shixuan Liu, Wenhao Wang, Changling Zhou, Chao Zhang, Jiandong Jin, Cheng Zhu. "A Unified Modeling Framework for Automated Penetration Testing", Computers & Security, 2025 Crossref	9 words — < 1%
89	corpora.tika.apache.org Internet	9 words — < 1%
90	decisionsciences.org Internet	9 words — < 1%
91	docs-library.unoda.org Internet	9 words — < 1%
92	dspace.vgtu.lt Internet	9 words — < 1%
93	editors.cis-india.org Internet	9 words — < 1%
94	idoc.pub Internet	9 words — < 1%
95	ijsate.com Internet	9 words — < 1%
96	legalinstruments.oecd.org Internet	

		9 words — < 1%
97	libweb.kpfu.ru Internet	9 words — < 1%
98	mediastorage.o-ran.org Internet	9 words — < 1%
99	multilingualbakbis.intgovforum.org Internet	9 words — < 1%
100	onlinelibrary.wiley.com Internet	9 words — < 1%
101	pariscall.international Internet	9 words — < 1%
102	pure.aber.ac.uk Internet	9 words — < 1%
103	researchmethodscommunity.sagepub.com Internet	9 words — < 1%
104	trepo.tuni.fi Internet	9 words — < 1%
105	www.bicg.eu Internet	9 words — < 1%
106	www.first.org Internet	9 words — < 1%
107	"Biologically Inspired Cognitive Architectures 2024", Springer Science and Business Media LLC, 2024 Crossref	8 words — < 1%

108 "Economics of Information Security and Privacy", Springer Science and Business Media LLC, 2010 — 8 words — < 1%
Crossref

109 "Risks and Security of Internet and Systems", Springer Science and Business Media LLC, 2025 — 8 words — < 1%
Crossref

110 Advances in Information Security, 2015. — 8 words — < 1%
Crossref

111 Andrew T. H. Tan. "The Global Arms Trade - A handbook", Routledge, 2014 — 8 words — < 1%
Publications

112 Arpita Nath Boruah, Mrinal Goswami, Manoj Kumar, Octavio Loyola-González. "Embedded Artificial Intelligence - Real-Life Applications and Case Studies", CRC Press, 2025 — 8 words — < 1%
Publications

113 Bill Jung, Yan Li, Tamir Bechor. "CAVP: A context-aware vulnerability prioritization model", Computers & Security, 2022 — 8 words — < 1%
Crossref

114 Brandon Wang, Xiaoye Li, Leandro P. de Aguiar, Daniel S. Menasche, Zubair Shafiq. "Characterizing and Modeling Patching Practices of Industrial Control Systems", Proceedings of the ACM on Measurement and Analysis of Computing Systems, 2017 — 8 words — < 1%
Crossref

115 Ikechukwu Oranekwu, Lavanya Elluri, Roberto Yus, Anantaa Kotal. "Scalable automation for IoT cyberSecurity compliance: Ontology-driven reasoning for real-time assessment", Computers & Security, 2026 — 8 words — < 1%
Crossref

116 Kamil Baraniuk, Piotr Marszałek. "The potential of Cyber Threat Intelligence analytical frameworks in research on information operations and influence operations", Przegląd Bezpieczeństwa Wewnętrznego, 2024
Crossref
8 words — < 1%

117 Kamor Kareem, Nitin Naik, Paul Jenkins, Paul Grace, Jingping Song. "Chapter 35 Understanding the Defence of Operational Technology (OT) Systems: A Comparison of Lockheed Martin's Cyber Kill Chain, MITRE ATT&CK Framework, and Diamond Model", Springer Science and Business Media LLC, 2024
Crossref
8 words — < 1%

118 LexisNexis
Publications
8 words — < 1%

119 Mohammed Almukaynizi, Alexander Grimm, Eric Nunes, Jana Shakarian, Paulo Shakarian. "Predicting Cyber Threats through Hacker Social Networks in Darkweb and Deepweb Forums", Proceedings of the 2017 International Conference of The Computational Social Science Society of the Americas on - CSS 2017, 2017
Crossref
8 words — < 1%

120 Mohiuddin Ahmed, Nazim Choudhury. "Cybersecurity for Internet of Health Things", CRC Press, 2025
Publications
8 words — < 1%

121 Mohiuddin Ahmed. "Ransomware Evolution", CRC Press, 2024
Publications
8 words — < 1%

122 Mohsin Raza, Ayesha Imtiaz, Umar Shoaib. "A Review on Security Issues and their Impact on Hybrid Cloud Computing Environment", International Journal of Advanced Computer Science and Applications, 2019
8 words — < 1%

Crossref

123 Netina Tan, Yuko Kasuya. "Routledge Handbook of Autocratization in Southeast Asia", Routledge, 2025 — 8 words — < 1%
Publications

124 Prasasthy Balasubramanian, Sonali Liyana, Hamsini Sankaran, Shambavi Sivaramakrishnan et al. "Generative AI for cyber threat intelligence: applications, challenges, and analysis of real-world case studies", Artificial Intelligence Review, 2025 — 8 words — < 1%
Crossref

125 Shannon Chance, Tom Børsen, Diana Adela Martin, Roland Tormey, Thomas Taro Lennerfors, Gunter Bombaerts. "The Routledge International Handbook of Engineering Ethics Education", Routledge, 2024 — 8 words — < 1%
Publications

126 Sobhy Abdelkader, Jeremiah Amissah, Sammy Kinga, Geofrey Mugerwa et al. "Securing Modern Power Systems: Implementing Comprehensive Strategies to Enhance Resilience and Reliability Against Cyber-Attacks", Results in Engineering, 2024 — 8 words — < 1%
Crossref

127 Tomer Simon, Ben Haklai. "Data Disaster – A New Threat for Sovereignty", International Journal of Disaster Risk Reduction, 2025 — 8 words — < 1%
Crossref

128 Volha Bryl. "Evaluating procedural alternatives: a case study in e-voting", Electronic Government an International Journal, 2009 — 8 words — < 1%
Crossref

129 apps.dtic.mil — 8 words — < 1%
Internet

130	dev.diplomacy.edu Internet	8 words — < 1%
131	doria.fi Internet	8 words — < 1%
132	eajournals.org Internet	8 words — < 1%
133	eprints.soton.ac.uk Internet	8 words — < 1%
134	etheses.whiterose.ac.uk Internet	8 words — < 1%
135	lists.oasis-open.org Internet	8 words — < 1%
136	media.defense.gov Internet	8 words — < 1%
137	mgaleg.maryland.gov Internet	8 words — < 1%
138	papers.academic-conferences.org Internet	8 words — < 1%
139	repository.kinnaird.edu.pk Internet	8 words — < 1%
140	securityanddefence.pl Internet	8 words — < 1%
141	static.rusi.org Internet	8 words — < 1%
142	web.mit.edu Internet	8 words — < 1%

143	www.enisa.europa.eu Internet	8 words — < 1%
144	www.intgovforum.org Internet	8 words — < 1%
145	www.opentext.com Internet	8 words — < 1%
146	www.picussecurity.com Internet	8 words — < 1%
147	www.thomasalspaugh.org Internet	8 words — < 1%
148	"Information Technology for Peace and Security", Springer Science and Business Media LLC, 2019 Crossref	7 words — < 1%
149	"Information and Communications Security", Springer Science and Business Media LLC, 2026 Crossref	7 words — < 1%
150	Maysa Sinan, Mojtaba Shahin, Iqbal Gondal. "Implementing and integrating security controls: A practitioners' perspective", Computers & Security, 2025 Crossref	7 words — < 1%
151	Uche Onyekpe, Vasile Palade, M. Arif Wani. "Recent Advances in Deep Learning Applications - New Techniques and Practical Examples", CRC Press, 2025 Publications	7 words — < 1%
152	repository.derby.ac.uk Internet	7 words — < 1%
153	"Information Fusion for Cyber-Security Analytics", Springer Science and Business Media LLC, 2017	6 words — < 1%

Crossref

154 "List of Icest 2023 Reviewers", 2023 58th International Scientific Conference on Information, Communication and Energy Systems and Technologies (ICEST), 2023 — 6 words — < 1%
Crossref

155 "NL ARMS Netherlands Annual Review of Military Studies 2020", Springer Science and Business Media LLC, 2021 — 6 words — < 1%
Crossref

156 "Network and System Security", Springer Science and Business Media LLC, 2020 — 6 words — < 1%
Crossref

157 Emil Larsson. "Collateral Damage from Offensive Cyber Operations—A Systematic Literature Review", Journal of Cybersecurity and Privacy, 2025 — 6 words — < 1%
Crossref

158 Lev Topor. "Cyber Sovereignty", Springer Science and Business Media LLC, 2024 — 6 words — < 1%
Crossref

159 Rafiq Ahmad Khan, Siffat Ullah Khan, Musaad Alzahrani, Muhammad Ilyas. "Security Assurance Model of Software Development for Global Software Development Vendors", IEEE Access, 2022 — 6 words — < 1%
Crossref

160 Ting-Han Chen, Jeroen van der Ham-de Vos. "Vulnerability Disclosure or Notification? Best Practices for Reaching Stakeholders at Scale", Digital Threats: Research and Practice, 2025 — 6 words — < 1%
Crossref

161 Vikas Khullar, Vrajesh Sharma, Mohit Angurala, Nipun Chhabra. "Artificial Intelligence and Society — 6 words — < 1%

5.0 - Issues, Opportunities, and Challenges", CRC Press, 2024
Publications

162	www.itu.int Internet	6 words — < 1%
163	archive.ceciis.foi.hr Internet	4 words — < 1%

EXCLUDE QUOTES	ON	EXCLUDE SOURCES	OFF
EXCLUDE BIBLIOGRAPHY	ON	EXCLUDE MATCHES	OFF