Review of Human Decision-making during Computer Security Incident Analysis

JONATHAN M. SPRING and PHYLLIS ILLARI, University College London

We review practical advice on decision-making during computer security incident response. Scope includes standards from the IETF, ISO, FIRST, and the US intelligence community. To focus on human decision-making, the scope is the evidence collection, analysis, and reporting phases of response, which includes human decision-making within and connecting these phases. The results indicate both strengths and gaps. A strength is available advice on how to accomplish many specific tasks. However, there is little guidance on how to prioritize tasks in limited time or how to interpret, generalize, and convincingly report results. Future work should focus on these gaps in explication and specification of decision-making during incident analysis.

CCS Concepts: • **Applied computing** → **Investigation techniques**; **Evidence collection, storage and analysis**; • **Security and privacy** → Security requirements;

Additional Key Words and Phrases: Security standards, sociology of science, cybersecurity

ACM Reference format:
Jonathan M. Spring and Phyllis Illari. 2021. Review of Human Decision-making during Computer Security Incident Analysis. *Digit. Threat.: Res. Pract.* 2, 2, Article 11 (April 2021), 47 pages.
https://doi.org/10.1145/3427787

1 INTRODUCTION

The purpose of this literature review is to identify what is publicly available concerning the structure of human decision-making during computer-security **Computer Security Incident Response (CSIR)** (shortened as Computer Security Incident Response (CSIR) if the usage is unambiguous) and identify the existing research programs that underpin the study of each element of the decision-making structure. The utility of such a literature review is to identify opportunities to expand research programs, as well as to identify loci of research programs best strengthened by interdisciplinary work or requiring interfield theories.

Incident response is an attractive topic, because it anchors the whole field of information security. When information security is not responding to an incident, it is either preparing for one or learning from past CSIRs.[1]

[1] Prevention of future incidents may fruitfully be discussed as independent from CSIR, but in practice prevention techniques are almost all adapted from the lessons learned after responding to an incident.

Supported by University College London's Overseas Research Scholarship and Graduate Research Scholarship. orcid.com/0000-0001-9356-219X.

Authors' addresses: J. M. Spring and P. Illari, University College London, Gower Street, London, WC1E 6BT; emails: jspring@cs.ucl.ac.uk; phyllis.illari@ucl.ac.uk.

2576-5337/2021/04-ART11
https://doi.org/10.1145/3427787

Preparation and learning are each complex and independent. However, their crux is CSIR. We restrict our focus to human-centric decision-making as opposed to automated or machine learning techniques. Automation is clearly crucial to CSIR, but deciding what automation to build and use remains a human decision, as does how to interpret and act on results. Yet, we suspect the structure of human decision-making is under-represented in available literature.

Cybersecurity[2] is important, yet breaches are usually detected months after they occur [183]. Prevention and detection are not always possible, so CSIR is an important part of cybersecurity. Successful guidance on how incident responders should make decisions during analysis would improve CSIR, and thereby all of information security. The research question we address with this review is:

RQ What are the gaps in public advice about how CSIR professionals should make, prioritize, and evaluate decisions during incident analysis?

This question is deceptively simple. It is easiest to illustrate what we seek in contrast to what we suspect is available. In summary, there is ample advice on what questions to ask and how to answer a given specific question, abundant tools fitted for well-defined narrow tasks, but little on less well-defined human decisions on issues such as how to prioritize questions and tasks, or interpret results. Without such guidance, an array of questions and tools falls far short of a satisfactory process by which incident responders generate and prioritize questions to ask, answer these questions under time constraints, and use these answers to make decisions. Therefore, we identify a significant lacuna in the existing literature on decision-making and evidence evaluation during CSIR.

The immediate challenge of this literature review is an explosion of scope. Scope can easily become too broad both because the topic is of broad application with many sub-parts and because the academic literature is not the only relevant source. Practitioners are an additional necessary source if the review is to adequately capture the state-of-the-art. Further, information security practitioners do not publish according to predictable norms, conveniently organized for input to an academic literature review. This means we have to design an appropriately flexible but tractable search and appraisal strategy.

Section 2 is therefore more detailed than might otherwise be expected. Section 2.1 defines CSIR and the review's scope within it, and Section 2.2 explains the relevant publication venues. With these definitions in hand, Section 3 is in position to demonstrate that no other literature review has covered this scope. Section 4 then sets about defining our search strategy (Section 4.1) and method for appraising documents (Section 4.2). The method mainly involves search for relevant standards documents and then mining them for citations. Section 4.3 presents the results of this first round—the relevant standards documents—and Section 4.4 presents the results of the citation mining. Section 5 presents our analysis of all the documents, organizes the relevant ones, and identifies three important gaps in the literature. Section 6 suggests how we might begin to fill these gaps. The three gaps we will identify in public advice to CSIR analysts are:

- Strategic selection of tactics, that is, which analysis heuristic or technical tool to employ in a particular situation and why
- When the investigator is justified in generalizing; that is, making a stronger, broader claim from singular pieces of evidence

[2] The field may variably be called information security, computer security, cybersecurity, or even the laborious "information and computer technology resiliency management." "Computer security" is an older term, and is canonized in terms such as Computer Security Incident Response Team (CSIRT). "Information security" emphasizes that human users and the physical world are part of the system under study, alongside computers. "Cybersecurity," etymologically at least, adds the social spaces that humans create using the computers to the socio-technical system under study [145]. We endeavor to use the same term as the source material, where possible. When making our own claims, we favor "cybersecurity" as the most expansive term.

- What information to report and how to communicate it to convince someone that the investigator should be believed

2 SCOPE

This section defines the scope of the literature review in two distinct aspects: the definition of the topic and the publication venues. Section 2.1 explains restricting the definition of CSIR to three subtasks during investigation: evidence collection, analysis, and reporting. Section 2.2 explains restricting the publication venues to relevant international standards and academic literature that is referenced therein.

As far as possible, we will use standard definitions for terms and prefer global, consensus, freely available definitions: in order of preference, the **Internet Engineering Task Force (IETF)**, **International Organization for Standardization (ISO)**, and **Forum of Incident Response and Security Teams (FIRST)**. This ordering is based on the extent of consensus (the Internet Engineering Task Force (IETF) membership is broader than Forum of Incident Response and Security Teams (FIRST)) and the openness of the definitions. Otherwise, the choice of established definitions for jargon is primarily for clarity, and to compress the discussion; we assume familiarity with the terms in the IETF Internet Security Glossary [156].

Section 3 provides evidence that the academic literature does not systematically cover our scope. We show this by a search through two common sources of literature reviews, **Association for Computer Machinery (ACM)** Computing Surveys and the **Institute of Electrical and Electronic Engineers (IEEE)** Security and Privacy "systematization of knowledge" papers. To summarize, no relevant surveys have been published on human decision-making during CSIR. While the reason why is unclear, a contributing factor may be the difficulty created by the natural secrecy of practitioners.[3]

A further reason we focus on standards as the starting point is simply to prevent an explosion of documents to review. A cursory Google Scholar search for "computer security incident response" and "digital forensic investigation" each return tens of thousands of results. Alternatively, searches in the Association for Computing Machinery (ACM) Guide to Computing Literature and Institute of Electrical and Electronic Engineers (IEEE) Xplore databases for "computer security incident response" return 23 and 24 results, respectively (searches on August 6, 2019). Many of these results are obscure, and for those that are not it is challenging to evaluate their operational impact. Standards are a better space to review, because, while some standards may be used more than others, the remit and authority of standards is explicit.

2.1 Scope—Topic

The first task is to explain what falls under CSIR and what it excludes. CSIR is a subspecies of business continuity planning or continuity of operations. In turn, continuity planning may be a response to man-made events (such as a military invasion) or natural events (such as a hurricane), and either physical or digital events. CSIR only includes response to primarily digital security incidents, where a security incident is something "contrary to system policy" [156]. Thus, accidents of all kinds are out of scope; though distinguishing apparent accidents from malicious acts is included. Intentional physical destruction of computing resources is also excluded from CSIR [23], but physical theft of a device holding cryptographic key material is in scope. CSIR is an activity that either Computer Security Incident Response Teams (CSIRTs) or **Product Security Incident Response Teams (PSIRTs)** may perform and CSIR may or may not require multi-party coordination; although CSIR adjusts and changes in these diverse contexts, we aim for the common themes.

[3] The comments from one anonymous reviewer are also instructive: "I am regularly attending meetings of practitioners [since 1993 and I do not feel] that secrecy is the issue. Certainly details of attackers' modus operandi have been protected. [B]ut almost always practitioners are not reasoning about WHY or SCIENCE, but HOW they solved the problem and what they learned about it. [B]ut again[,] on the technical level and not on the process level. So while clearly you need to have a process to cover evidence collection and attribution ...there is simply no body of knowledge around [it] in the incident response community."

Narrowing the focus further, CSIR is a task within incident management.[4] CERT/CC definitions of incident management [2, 132] locate CSIR as independent from activities such as preparation, improving defenses, training, financing, and lessons learned. Mundie et al. [132] surveys practices including those by CERT/CC and ISO; six tasks are included as part of CSIR: monitoring, detection, evidence collection, analysis, reporting, and recovery. For a multi-party incident, we might add coordination [16] as an overlay connecting multiple analysts conducting CSIR; we restrict our scope to the analysis process of a single analyst.

These six tasks form the core topic of this survey of CSIR. However, the human-centric decisions that are elements of these six CSIR tasks vary in importance. Analysis, reporting, and recovery are almost wholly human-driven, and monitoring is almost wholly automated, while detection and evidence collection are a mixture. Where detection is automated, say, in an **intrusion detection system (IDS)**,[5] it is out of scope. Decisions about what detection rules to implement in an IDS are part of the preparation or improving defenses phases of incident management, as a result of lessons learned, and thus are also out of scope. Actual human intrusion detection is rare, and when it occurs is usually the result of analysis during CSIR to some other, automatically detected incident. Therefore, our focus on human-driven CSIR investigation excludes monitoring and detection.

The IETF [23, 156] and CERT/CC define neither "investigation" nor "forensics" in relation to the incident management process. ISO/IEC [90] places investigation as the centerpiece of incident management, where the principles of incident management are to "give guidance on the investigation of, and preparation to investigate, information security incidents" ISO/IEC [91, §;0]. In this way, ISO uses "investigation" as a near-synonym to "response" in the IETF and FIRST literature.

Our use of "incident" emphasizes that the investigation is oriented towards the violation of some policy, possibly but not necessarily a law. Thus, modelling or analyzing online crime is an investigation, and so is IT staff looking into a usage policy violation. Incident response or investigation is entwined with cybersecurity, because one essential aspect of a defense strategy is feedback from investigation to "preparation" and "protection" [2]. Detailed discussion of preparation and protection is placed out of scope, because the relationship is complex and our scope is large enough by discussing CSIR; however, "reporting" covers how an analyst should make this link to preparation and protection.

Incident response, per IETF and FIRST, explicitly includes remediation, but ISO [90] treats remediation and response as separate from investigation. In determining scope, we follow ISO and exclude remediation, which the other documents call recovery. Note that both sets of standards agree that clear reporting is the proper output of incident analysis, and any recovery follows reporting. However, it does seem clear that recovery follows a different decision process than analysis, and the two should be treated separately. Within the six tasks identified within CSIR, three are left in scope:

- evidence collection
- analysis
- reporting

[4] The term "incident management" does not appear in IETF documents consistently. Trammell [177] describes Incident Object Description Exchange Format (IODEF) [47] as a protocol for "exchange of incident management data," but the term "incident management" does not appear again in Trammell [177], and not once in Danyliw et al. [47]. ISO/IEC [91] defines "information security incident management" as "exercise of a consistent and effective approach to the handling of information security incidents." FIRST defines an "information security incident management" service area [16], and also FIRST [60] recommends the CERT® Coordination Center operated by Carnegie Mellon University (CERT/CC) documentation on incident management. Trammell and Danyliw both worked at CERT/CC, which may explain the informal reference in the IETF documents. The CERT/CC phases are consistent with the ISO/IEC [91] phases of plan and prepare; detection and reporting; assessment and decision; responses; and lessons learned. We prefer the CERT/CC definitions, as they are public (vice the International Organization for Standardization (ISO) standards) and recommended by FIRST (thus in scope of using global, consensus-driven definitions).

[5] Shirey [156] refers to Bace and Mell [10] for intrusion detection system (IDS) details, which has been superseded by Scarfone and Mell [153].

These three seem too tightly coupled to separate and are described consistently across the international standards organizations, and all three involve human decision-making.

For each of these three topics, our concern is primarily with how an individual analyst or team makes decisions during these three phases. What tool or what language the analyst or investigator uses to make these choices is not germane and is out of scope. This is not a review of available security tools, platforms, or data exchange formats. Our goal is to survey how analysts enumerate options, including of tasks and tools, evaluate choices, including what questions to prioritize, generalize results, and justify these steps.

Our scope topic does not directly address "digital forensic investigation." The connotations of this term are influenced by the fact it is the term of art within the law enforcement community. A coarse distinction between CSIR and digital forensic investigation is that both are terms for analyzing and responding to incidents on computers; CSIR is for the context of security policy violations, whereas digital forensic investigation is for the context of legal, especially criminal, investigations. These different contexts create important process differences. But the decision-making process in both types of analysis are related. Our selection of publication venues focuses on CSIR standards, but many of them cite digital forensic investigation best practice documents. Thus, our review will map out the aspects of decision-making in digital forensics that have influenced CSIR. However, we do not make an independent survey of digital forensics standards, because we are seeking to understand what has been adopted as standard practice within CSIR, not what else digital forensics might add in the future.

2.2 Scope—Publication Venues

As CSIR and investigation includes professional and business aspects, viable sources on CSIR practices are not limited to academic sources. As Spring et al. [165] documents, the science of security is an unsettled area of research rather than an area with anything like standards. In fact, traditional academic publication venues contain little if anything about day-to-day CSIR practices; academics do not do CSIR themselves. Sundaramurthy et al. [171] seems to mark the first anthropological study of a CSIRT[6] members and their attitudes, but this literature is not about the actual process of CSIR; that is covered in the professional literature. Historiographic study of CSIRTs is even more recent: "CSIRTs have been around for more than 30 years, but little is known about their beginnings" [158]. And standards are key to understanding of CSIR: "Trust in infrastructure, we suggest, requires not only trusted standards, but also trust in the actors and organizations that implement and maintain those standards" [158, p. 174]. The search venues will be the IETF, ISO,[7] FIRST, and documents understood to represent the US **intelligence community (IC)**.

Therefore, to understand current CSIR practices the scope of the review is internationally relevant standards and whatever literature is referenced therein. The history of standards as its own industry is complex in its own right [166]. The Internet and IT standards are formed by heterogeneous processes involving a wide variety of actors [135]. Security-relevant standards are beginning to be seen as having their own unique requirements, distinct from IT standards generally [106]. However, it is a separate project to analyze how CSIR standards have come to be. The standards in this review are taken as-is, with the understanding that any interpretations should be made cautiously because the standards may not cleanly fit in to existing studies of how and why other IT standards are created. More than other IT standards, CSIR standards are likely a codification of tacit practitioner knowledge [133].

The scope is not restricted to the traditional academic venues, as we wish to focus on what incident responders actually do. Ideally this would take the form of first-hand accounts; however, cybersecurity is a sensitive topic.

[6] CSIRT is the general term, and will be used unless referring to a specific organization.

[7] Although ISO standards are only available for a fee, the terms and definitions as used in the relevant standards (the 27000 series) are freely available.

Chatham House Rules[8] are common, practitioners routinely request information dissemination restrictions with **Traffic Light Protocol (TLP)**, and **non-disclosure agreements (NDAs)** abound in the discipline. These norms within the security community further frustrate the usual academic publication expectations, as it would be impossible to evaluate the selection bias or outright deception within studies. The CSIR standards at least form an honest baseline of what is expected of a competent practitioner, and this review applies to competent practitioners, where competent is defined by consensus among practitioners and codified in the standards. However, we do not empirically address the extent to which competence is common. Due to the community norms of secrecy (documented by Sundaramurthy et al. [171]), a comprehensive evaluation is impractical.

The scope of publication venues is limited to ISO, IETF, FIRST, and the US IC. This choice is based on what organizations are relevant in influencing or describing international CSIR practices, which in turn is due to the history of the governance of the Internet. We mitigate potential over-restriction of focus by including any documents cited by standards publications. Our reasoning for selecting these organizations specifically is as follows:

ISO and the **International Telecommunications Union (ICU)** are the authoritative technology standards makers [135, p. 11]. The US federal government plays a dominant role in Internet development and standards, through the original **Advanced Research Projects Agency (ARPA)** development under the US **Department of Defense (DoD)** and subsequent stewardship under the Department of Commerce.[9]

ISO is *de dicto* where one looks for international standards. Each nation-state is allowed to have one member in ISO, namely, the official national standards body representing all the industry-based bodies in each country. It is a federation of federations, representing a multitude of industries. ISO standardizes things like the two-letter country codes (which have been adopted as **Domain Name System (DNS)** top-level domains), paper sizes, and credit cards. The International Telecommunications Union, an agency of the UN (ITU) and their CIRT program[10] seems promising in name; however, their website publishes little besides an events list. It appears that content is provided by FIRST members, companies, or other consultancies; the ITU does not produce its own CSIR materials or standards. This leaves only ISO in scope of the potential authoritative international standards bodies.

However, the IETF is the *de facto* place to go for international Internet standards because, for all intents and purposes, its standards are the Internet. The IETF "doesn't recognize kings—only running code" and creates more pragmatic, open (freely available) standards [135, p. 12]. Open standards happen to have won out on the Internet; IETF standards like **Transmission Control Protocol/Internet Protocol (TCP/IP)**, Domain Name System (DNS), and **Border Gateway Protocol (BGP)** underpin every Internet connection. For a background history of how the precursor to the IETF came to this dominant role, see Hafner and Lyon [72]. The other main open-standards body is the **World Wide Web Consortium (W3C)**, which standardizes **Hypertext Transfer Protocol (HTTP)** and **Extensible Markup Language (XML)**, for example. World Wide Web Consortium (W3C) stays narrowly focused on web standards, and although this includes important web security considerations, W3C does not work on incident management, so we mark the group as out of scope.

FIRST is not part of this longer **information and communications technology (ICT)** standards history. It was formed in 1990 specifically to coordinate among and represent the interests of CSIRTs globally. FIRST's mission includes developing and sharing best practices, as well as creating and expanding CSIR teams [59]. FIRST is the one and only global organization representing those who do human-centric CSIR tasks. FIRST's work with **United Nations (UN)** agencies like the ITU also testifies to its global influence. It is naturally included as in-scope.

[8] Chatham House Rules indicates a situation in which the information or content of a meeting, discussion, or presentation may be disclosed but the source of the information may not be identified, implicitly or explicitly. This request is made by the speaker prior to disclosing the information.

[9] Two important sub-parts of Commerce are Internet governance by the National Telecommunications and Information Administration and standards by National Institute of Standards and Technology, part of the US Department of Commerce (NIST).

[10] http://www.itu.int/en/ITU-D/Cybersecurity/Pages/Organizational-Structures.aspx.

There are three organizations one might consider naturally in scope that are excluded. These are **EU Agency for Network and Information Security (ENISA)**, National Institute of Standards and Technology, part of the US Department of Commerce (NIST), and the United States of America (US) US Department of Defense (DoD). However, within the gray area between NIST and the US intelligence community, we identify a fourth set of *de facto* standards.

EU Agency for Network and Information Security (ENISA) is specifically focused on CSIRTs and information security. It is focused on coordination between member CSIRTs and advice on **European Union (EU)** policies, which does not make a good fit with our scope of the analysis process of a single CSIR analyst. The European Union (EU) publishes an independent evaluation of ENISA's activities.[11] A key document interfacing with EU standards is an extended definition of the term "cybersecurity" and what EU work is done related to it [22]. EU directive 2016/1148 increased ENISA's statutory powers when it came into effect in November 2018, but ENISA's focus on coordination and policy remains the same. ENISA has CSIRT services divided into "reactive" and "proactive." There are four publications on reactive services, including "incident response," between November 1, 2013, and September 1, 2016; there are none between then and June 1, 2020.[12] The publications are about tools for identifying and processing "actionable information" and have little to do with our topic as defined in Section 2.1, so we leave ENISA out of scope.

NIST is a difficult organization to place in or out of scope. It is part of the Department of Commerce, and so has loose ties to the remaining Internet stewardship invested in the National Telecommunications and Information Administration. Strictly, NIST merely sets policies for how the US federal civilian government secures its IT infrastructure and responds to incidents [41]. This **Federal Information Security Management Act (FISMA)** policy responsibility is a part of NIST's larger role of "advancing measurement science, standards, and technology in ways that enhance economic security and improve our quality of life" [178]. Through this role, NIST standardized AES, which is the *de facto* global standard encryption algorithm. NIST documents and standards are also cited by the IETF, ISO, and FIRST, elevating certain NIST work from a national to international status. We shall consider NIST generally out of scope; however, many NIST publications will be considered as works cited by the international standards organizations.

There are two US federal government units that do not fall under NIST's authority—the DoD and the **Central Intelligence Agency (CIA)**. These two organizations have not published CSIR standards as openly as NIST or these other standards organizations. DoD has some history of driving information security norms, so it is plausible to ask about its influence on CSIR. For example, a DoD document that has become canonical within the security community is the Orange Book [20].

The questions the DoD and its sub-agency the **National Security Agency (NSA)** have raised around whether cybersecurity is, broadly, a science (see, e.g., Galison [62], Katz [100], MITRE Corporation [126]) could inform evidence evaluation in incidence response, because evaluating evidence properly is a primary scientific activity. While these DoD projects ask the right questions about science to help with CSIR, they have generally concluded security is not (yet) a science, and so there is little advice. Spring et al. [165] argues this conclusion is ill-founded and excessively pessimistic. However, the relevant point for this review is that the science of security literature does not advise CSIR.

While the main part of the DoD does not publish adequate documents, the intelligence community aspects of the US federal government do. The DoD and Central Intelligence Agency (US) (CIA) are generally not forthcoming with more conventional descriptions of their CSIR practice. However, given that NIST is not authoritative over the intelligence community (IC), one would expect them to develop their own standard practices. Documents related to the practice of the US IC are occasionally published, with IC attribution either explicit or implicit.

[11] "Annual ex-post evaluation of ENISA activities" https://www.enisa.europa.eu/about-enisa/annual-ex-post-evaluation-of-enisa-activities.

[12] According to ENISA's search and filter functionality located at https://www.enisa.europa.eu/topics/csirt-cert-services/reactive-services?tab=publications.

Three such documents are relevant to evidence collection and analysis in incident investigation, forming what is essentially a *de facto* standard. The first is a textbook published by the CIA and used to train intelligence analysts [79] whose methods are applicable to CSIR. The second is a pair of documents, the kill chain model of computer network attacks [84] and the diamond model of intrusion analysis [27]. Unlike the textbook, these documents are not explicitly acknowledged as standard analysis methods within the defense and intelligence communities. However, the diamond model paper is published by the DoD publisher, the Defense Technical Information Center.[13] The diamond model builds on the kill chain. Given that Lockheed Martin, a US defense contractor, published the kill chain, it seems the papers are from overlapping communities. Although it is tenuous to term three documents a "standard," it is clear from the content that they come from a practitioner community and are one of the clearest publicly available expressions of intrusion investigation methods. Therefore, they are clearly in scope for discussion.

The US intelligence agencies exercise out-sized international influence. The US is part of an intelligence sharing alliance known as the five eyes, which includes Australia, Canada, New Zealand, and the United Kingdom. As the biggest partner in this alliance by far, what the US intelligence practitioners do is probably accommodated, if not directly copied, by the other countries' services.

US military influence goes beyond even the five eyes. The **North Atlantic Treaty Organization (NATO)** is the biggest alliance the US leads, with 28 other countries. North Atlantic Treaty Organization (NATO) intelligence is also presumably influenced by five eyes, as Canada and the UK also play a big role. The US tends to supply logistics and intelligence support in its alliances, so intelligence standards are likely to influence allies. Other locations that cooperate extensively with the US include Israel, South Korea, Japan, and the Philippines. By virtue of these alliances, it is reasonable to assume that intelligence professionals in all these places are relatively closely aligned with US intelligence standards. These alliances end up including most of the global military and intelligence spending. Essentially only China and Russia are excluded, and the two of them account for 15%–20% of global military spending. Thus, although there are rather few IC documents, and they are focused on the US, they should provide information about how a large swath of such practitioners make decisions.

In summary, this review will include the IETF, ISO, FIRST, and available intelligence community documents as in-scope publication venues for incident investigation standards of practice for evidence collection, analysis, and reporting. The review will exclude the ITU, W3C, ENISA, and US federal civilian government departments and agencies as out of scope due to either limited applicable content or limited jurisdiction. The most borderline organization is NIST, which occasionally has standards canonized by the in-scope venues; the review will only include those NIST standards cited or adopted explicitly by the four in-scope venues. Section 4.1 describes the method for determining which standards are relevant within these venues.

3 NOVELTY

This section demonstrates that our intended scope, as defined in Section 2, has not been previously surveyed, by providing a brief structured survey of the survey literature. Evidence is provided by the lack of related surveys in two academic venues: IEEE Security and Privacy **Systematization of Knowledge (SoK)** papers and **ACM Computing Surveys (CSUR)** journal. Here, all 42 extant SoK papers (as of August 1, 2019) are appraised for relevance, while for ACM CSUR, we apply a keyword search to the corpus.

IEEE S&P has published 42 SoK papers since the venue initiated the SoK publication option in 2010 through the 2019 conference. We make an exhaustive evaluation of relevance based on title and abstracts. Our basic relevance criterion in this case is if the SoK is about designing or evaluating investigations of maliciousness. Of the 42 SoK papers, only Herley and van Oorschot [78] and Rossow et al. [150] are applicable to our project. Prior work explains why each of these is insufficient for our purposes. Spring et al. [165] addresses the shortcomings

[13] See http://www.dtic.mil/docs/citations/ADA586960.

Table 1. Potentially Relevant Literature Reviews from ACM CSUR

	Found in search #								Criteria		
Document	1	2	3	4	5	6	7	8	1	2	3
Laube and Böhme [109]	■		■	■	■	■	■	■	✓	✗	✗
Homoliak et al. [80]*		■							✓	✗	✓
Li et al. [114]			■						✗	✗	✓
Li et al. [113]			■						✓	✗	✗
Barmpatsalou et al. [12]			■			■			✓	✓	✗
Islam et al. [86]			■	■	■	■	■	■	✓	✗	✓
Jiang et al. [93]				■					✗	✗	✗
Cho et al. [38]				■	■		■	■	✗	✗	✗
Liu et al. [116]				■	■		■	■	✗	✓	✗
Botacin et al. [19]				■	■		■	■	✓	✓	✗
Ramaki et al. [146]				■	■	■			✗	✓	✗
Jhaveri et al. [92]					■		■	■	✓	✗	✓
Pendleton et al. [142]					■		■	■	✗	✗	✗
Khan et al. [103]					■		■	■	✓	✓	✗
Laszka et al. [108]						■	■	■	✗	✓	✓
Kalgutkar et al. [99]						■	■	■	✗	✗	✗
Biddle et al. [17]							■	■	✗	✗	✗
Milenkoski et al. [123]							■	■	✗	✓	✗
Tang et al. [174]							■		✗	✗	✗
Meng et al. [121]							■	■	✗	✗	✗
Calzavara et al. [28]							■	■	✗	✓	✗
Labati et al. [107]							■	■	✗	✓	✗
Ye et al. [189]							■	■	✗	✓	✗
Edwards et al. [51]*							■	■	✓	✓	✓
Avancha et al. [8]							■	■	✗	✗	✗
Roy et al. [152]							■	■	✗	✓	✗
Chandola et al. [34]							■	■	✓	✓	✗
Pearce et al. [141]								■	✓	✓	✗
Peng et al. [143]								■	✗	✓	✗
Younan et al. [190]								■	✗	✓	✗
Egele et al. [52]*								■	✓	✓	✗

The three relevance criteria are (1) relates to forensics rather than prediction; (2) technical, investigative focus; (3) useful level of abstraction of CSIR. Papers with an asterisk (*) are discussed in more detail in the text.

in Herley and van Oorschot [78]. Hatleback and Spring [74] expands and generalizes the good work of Rossow et al. [150]. None of the SoK papers systematize knowledge of CSIR, investigation, or analysis.

Our CSUR keyword search uses Google Scholar, limiting to publications in "ACM Computing Surveys" between January 1, 2007, and August 1, 2019. We use all the keywords used in the main study, as described in Section 4.1. However, CSUR is a sufficiently different venue from our intended scope that we sometimes find different keywords more useful. The surveys returned by the following search terms are included in Table 1. Quotes are applied to the search as listed.

(1) "computer security CSIR"
(2) "incident investigation"

(3) "incident management"
(4) "computer security" & "evidence collection"
(5) "CSIR" & analysis
(6) "security incident" & "investigation"
(7) "computer security" & incident investigation
(8) "computer security" & incident analysis

Two search terms were tried on CSUR but returned too many clearly irrelevant results to be considered useful. Namely, {"computer security" & analysis} with 79 results and {"computer security" & reporting} with 25. Any relevant papers appear to be included in the 31 captured by Table 1.

These eight searches within CSUR return 31 unique results. As the search terms are expanded to include more general, related terms, we find a handful of possibly relevant results. To determine whether any of these surveys already adequately cover our topic of interest, we set out three relevance criteria. The survey must:

(1) relate to reconstructing past events (i.e., forensics) rather than prediction;
(2) focus on the technical- and knowledge-based decisions and processes, rather than management processes;
(3) use our target level of abstraction to discuss the problem of incident response, investigation, or analysis (human decisions during the process), rather than tool development, without being so abstract as to make implementation impractical.

These criteria are marked in Table 1, first based on each paper's abstract, although some papers deserved a look beyond their abstracts.

Edwards et al. [51] is the only survey that meets all three criteria, based on their abstract. However, their focus is quite different from our intended focus. They discuss automation of law-enforcement criminal investigation using computer science techniques. There may be overlap with computer-security CSIR, in that some subset of law enforcement cases involve criminal action against computers. However, the focus of Edwards et al. [51] is what Anderson et al. [6, p. 3] call, quoting the European Commission, "traditional forms of crime... committed over electronic communication networks and information systems." Incident response and investigation focuses on a different category, "crimes unique to electronic networks," as well as organizational policy violations that are not illegal under the relevant jurisdiction. Finally, Edwards et al. [51] focus on automation of police tasks, whereas our focus would be on the investigator's decision process in, among other things, choosing which automation techniques to use and how to evaluate the evidence they provide. These various differences make a clear case that our intended survey topic is sufficiently distinct from Edwards et al. [51].

Egele et al. [52, p. 1] aims to identify "techniques to assist human analysts in assessing ... whether a given sample deserves closer manual inspection." It is, in fact, a survey of software tools and their features, and does not discuss how an analyst should use them.

Homoliak et al. [80] is broadly about insider threat. This topic is a subset of computer security incidents, albeit with specific characteristics. However, Homoliak et al. [80] skirts our topics of interest; the authors discuss sets of case studies of analyzing insider threat cases, and a common structure for how insiders conduct attacks. For this first aspect, the authors state "[t]he majority of the research aimed at obtaining information about insider incidents was conducted by CERT..." (p. 16). We will address a more comprehensive analysis of advice from CERT/CC in the context of international CSIR norms put forward by FIRST. On the topic of incident structure, we will argue that the US IC has a better claim to the standard model, and Homoliak et al. [80] do not discuss the IC models.

Many papers in Table 1 meet the technical criterion (#2) and fail the other two criteria. This pattern tends to be about some specific subset of network defense—for example, making better passwords, intrusion detection systems, or web browser defenses. These tools are certainly used and evaluated as part of security management, and are important considerations. However, these details are tangential to making decisions during CSIR and investigation. As we suggested in Section 1, there are plenty of tools available.

However these reviews of the available survey literature demonstrate a lacuna; we lack a survey of CSIR and investigation practices. This omission matters. The outputs from incident analysts and CSIR inform and shape all other aspects of cybersecurity. Security research and security management requires, directly or indirectly, facts on which to build research or risk management plans. The practice of cybersecurity should work on the basis of evidence, and evaluate any other security infrastructure, plans, defenses, or policy in respect of that evidence. Incident analysis is perhaps the most important source of such evidence. The primary goal of cybersecurity is reducing the frequency and impact of incidents, and to do so understanding past incidents helps tremendously. However, it seems there is no systematic review of how this evidence should be collected, analyzed, and reported. One must understand these steps to properly interpret any such evidence. Therefore, although our topic is narrow, it has far-reaching impact on information security more generally.

4 SEARCH AND APPRAISAL

We have explained that the scope of our topic is restricted to evidence collection, analysis, and reporting in human-driven computer security CSIR. We further restrict our review to internationally recognized standards, to keep us as in touch as possible with actual professional practice without violating confidentiality around CSIR, which organizations often, justifiably, do not disclose in detail.

We have also explained the various challenges of designing an appropriate review strategy covering these questions. This section presents our search of the literature, appraisal of which of the resulting standards to include in our later synthesis, and further appraisal of references cited in the originally found standards documents. Our search and appraisal results can be found in Section 4.3, and referenced documents are discussed in Section 4.4. First, we explain our search and appraisal methodologies.

4.1 Methodology of Search Strategy

The major determining factor in our literature search strategy is the scope of publication venues, as we have justified in Section 2.2. Further searches had to be tailored to each venue.

Each of IETF,[14] FIRST,[15] and ISO[16] have dedicated web pages. For IETF and ISO, we use their site-based search engines that cover their respective corpora of standards, and use the following search terms there. Quotes are applied to the search as listed. The keywords employed are:

(1) "computer security incident response"
(2) "incident investigation"
(3) "incident management"
(4) "computer security" & "evidence collection"
(5) "computer security" & analysis
(6) "computer security" & reporting

We added or modified terms slightly to accommodate each search venue. The IETF RFC search tool does not accommodate mixing quoted phrases with other terms, so for terms 4, 5, and 6 the quotes were removed. We added the following terms to the IETF search:

- "incident response"

[14] https://www.rfc-editor.org/search/rfc_search.php.
[15] https://first.org/standards/.
[16] https://www.iso.org/standards.html.

We added the following terms to the ISO search, after it became apparent from searches 2 and 3 that the ISO documents do not use the term "computer security" but rather "information security":

- "information security" & "evidence collection"
- "information security" & analysis
- "information security" & reporting

FIRST has a smaller, more focused corpus of work. It only lists four standards on its "standards" webpage, so we exhaustively evaluate these. The FIRST website has three other headings under which it publishes or links to advice for CSIR professionals: Security Reference Index,[17] Best Practice Guide Library,[18] and Education Program (Services Framework).[19] These lists are explicitly "not meant to be a definitive list," but they do capture resources, best practices, and content that an incident analyst might be expected to know. The lists of FIRST-developed content are straightforward to search for the above keywords. Where the lists contain primarily links to other organization's websites or documents, we search the linked resources for the keywords but do not follow further links.

This approach was not possible with the IC documents, although, as we have explained, they need to be included in our scope. Due to the idiosyncratic nature of the IC publication and publicity processes, there is no sense in a keyword search strategy. Instead, we have arrived at the core documents we shall consider as the "standards" from this community via essentially the only viable route: We know these to be key documents under Chatham House rules.

All the documents returned by this search strategy were appraised using the methods of Section 4.2. Eleven standards documents pass our appraisal, as Section 4.3 documents. We then take a further search step and extract the references from those eleven documents. Section 4.4 presents the results of this search and appraisal step. Our cutoff date for inclusion in this review is that the document must be published by August 1, 2019, and we only include any cited documents that are publicly available (or, in the case of ISO, readily available).

4.2 Methodology of Appraisal Strategy

The purpose of the appraisal is to determine whether each document is within the scope of evidence collection, analysis, and reporting for CSIR. All the documents extracted from the references were appraised using the same methods to determine whether they are included in our review, independent of the document that cited it.

Our specific inclusion criteria for whether the content is in-scope are the following:

- Target audience as expressed by author includes security professionals
- Topic applies to one of the following parts of computer-security CSIR (or some clear synonym thereof)
  –evidence collection
  –analysis
  –reporting
- Topic is on investigator practice (rather than implementation of software or managerial considerations related to CSIRTs)
- Document is finalized (not a draft) and not explicitly superseded as of August 1, 2019
- Document is available in English

A document must satisfy all of these criteria to be included in the review.

Standards may be superseded or amended by future work, because they follow an orderly progression, and drafts are commonly published for public comment before being finalized. We exclude any standard superseded

[17] https://www.first.org/resources/guides/reference.
[18] https://www.first.org/resources/guides/.
[19] https://www.first.org/education/services-framework.

Table 2. IETF Database Search Results

	Criteria				
Document	1	2	3	4	5
RFC 2350 [23]	✓	✗	✓	✓	✓
RFC 3607 [110]	✓	✗	✗	✓	✓
RFC 5070 [47]	✓	✓	✓	✗	✓
RFC 6045 [128]	✓	✓	✓	✗	✓
RFC 6046 [130]	✓	✓	✓	✗	✓
RFC 6545 [129]	✓	✓	✓	✓	✓
RFC 6546 [176]	✓	✗	✓	✓	✓
RFC 7203 [173]	✓	✓	✓	✓	✓
RFC 7970 [46]	✓	✓	✓	✓	✓
RFC 8134 [85]	✓	✓	✓	✓	✓

The criteria are (1) target audience is security professionals; (2) topic in scope, per Section 2.1; (3) focus is investigator practices; (4) document finalized and not obsoleted as of August 1, 2019; (5) available in English.

as of August 1, 2019, and incorporate any amendments finalized by August 1, 2019. We note the existence of drafts on new topics, but exclude their content from the review. Restriction to English-language documents should not distort the review, although see discussion of this and other potential limitations in Appendix B.

4.3 Search and Appraisal

We present our results according to search venue in Sections 4.3.1 through 4.3.4. We will synthesize the results in Section 5.

4.3.1 IETF. Table 2 evaluates the results of our search procedure for IETF. Four documents pass appraisal, RFCs 6545, 7203, 7970, and 8134.

The IETF documents break down into two clear broad categories, **Best Current Practice (BCP)** 21 on expectations for computer security incident response [23], and all the others, which are to do with **Incident Object Description Exchange Formant (IODEF)**, its expansion, and usage.

As an expectations document, Best Current Practice, a series of documents published by IETF (BCP) 21 focuses primarily on the services and support a CSIRT should provide to its constituency, who that constituency should include, and so on. These considerations are vital to CSIRT operations; however they are not directly relevant to our question at hand.

The IODEF projects in particular feature CERT/CC staff heavily. Danyliw and Inacio during all their RFC authorship, and Trammell contributed heavily to **SiLK** (**System for Internet-level Knowledge)** while at CERT/CC before moving on. There may be an incidental division of labor between the IETF documents and the FIRST documents, since many CERT/CC staff were involved in both, but there is no documentation any division was explicitly intentional. In particular, the IODEF format focuses almost exclusively on technical issues of data exchange and reporting format. The softer considerations, of how to collect, evaluate, and analyze the data contained within IODEF are in the purview of FIRST.

Thus, as technical reporting formats are in our scope of reporting results, all RFCs related to IODEF are relevant. There do not appear to be any other IETF documents within our scope.

These IODEF documents may at first seem to be out of scope, as we specified that our scope is how investigators make decisions, not what tools or formats they use to document them. This topic recurs in Section 4.4.1. IODEF is essentially a language for talking about computer security incidents. However, because we are not interested in

Table 3. ISO Database Search Results

	Criteria				
Document	1	2	3	4	5
IEC 31010:2009	✓	✗	✓	✓	✓
ISO 13485:2003	✗	✗	✗	✗	✓
ISO/IEC 17799:2005	✓	✗	✗	✗	✓
ISO 22320:2018	✗	✗	✗	✓	✓
ISO 22319:2017	✗	✗	✗	✓	✓
ISO/IEC 27000:2009	✓	✗	✗	✗	✓
ISO/IEC 27001:2005	✓	✗	✗	✗	✓
ISO/IEC 27002:2005	✓	✗	✗	✗	✓
ISO/IEC 27004:2016	✓	✗	✗	✓	✓
ISO/IEC 27005:2011	✓	✗	✗	✓	✓
ISO/IEC 27006:2011	✗	✗	✗	✗	✓
ISO/IEC 27006:2015	✗	✗	✗	✓	✓
ISO/IEC 27033-1:2009	✓	✗	✗	✗	✓
ISO/IEC 27033-4:2014	✓	✗	✗	✓	✓
ISO/IEC 27035:2011	✓	✓	✓	✗	✓
ISO/IEC 27035-1:2016	✓	✓	✓	✓	✓
ISO/IEC 27035-2:2016	✓	✗	✗	✓	✓
ISO/IEC 27041:2015	✓	✓	✓	✓	✓
ISO/IEC 27043:2015	✓	✓	✓	✓	✓
ISO/IEC TR 18044:2004	✗	✓	✓	✗	✓
ISO/IEC TR 20004:2015	✓	✗	✗	✓	✓
ISO/NP TS 11633-1	✓	✗	✗	✗	✓
ISO/TR 11633-1:2009	✓	✗	✗	✓	✓
ISO/TR 11633-2:2009	✓	✗	✗	✓	✓
ISO/TS 19299:2015	✗	✗	✗	✓	✓
ISO/TS 22330:2018	✗	✗	✗	✓	✓

The criteria are (1) target audience is security professionals; (2) topic in scope, per Section 2.1; (3) focus is investigator practices; (4) document finalized and not obsoleted as of August 1, 2019; (5) available in English.

data formats, we are not interested in the language per se. IODEF is in scope because as a constructed language it makes judgments about what aspects of incidents are important, necessary, or possible to communicate. These judgments, at least implicitly, bear on what an investigator should choose to report. We therefore judge IODEF as in-scope. However, we stress that our intended scope is how to decide what information to report, not what language in which to report it. Therefore, data formats and languages for anything else remain out of scope.

4.3.2 ISO. Nine search terms return 34 total results, with 26 unique results displayed in Table 3. ISO 27035-1, 27041, and 27043 meet our appraisal criteria to carry through to the citation-harvesting and synthesis stage:

- *Information security incident management, Part 1: Principles of incident management* [91]
- *Guidance on assuring suitability and adequacy of incident investigative method* [88]
- *Incident investigation principles and processes* [90]

Table 4. FIRST Results Summary

	Criteria				
Document	1	2	3	4	5
Mundie et al. [132]	✓	✓	✓	✓	✓
Alberts et al. [2]	✓	✓	✓	✓	✓
OCTAVE [29]	✓	✗	✗	✓	✓
ENISA [54]	✗	✓	✓	✓	✓
Cormack [45]	✓	✗	✗	✓	✓
Gorzelak et al. [66]	✓	✓	✓	✓	✓
Cichonski et al. [41]	✓	✓	✓	✓	✓
ETSI [55]	✓	✗	✓	✓	✓
RFC 2350 [23]	✓	✗	✓	✓	✓
RFC 2196 [61, Section 5.4 only]	✓	✓	✓	✓	✓
RFC 2827 [58]	✓	✗	✗	✓	✓
RFC 2504 [71]	✗	✓	✓	✓	✓
Services Framework v2 [16]	✓	✓	✗	✗	✓

The criteria are (1) target audience is security professionals; (2) topic in scope, per Section 2.1; (3) focus is investigator practices; (4) document finalized as of August 1, 2019; (5) available in English.

4.3.3 FIRST. FIRST is the smallest body surveyed, and it is not primarily a standards organization but rather a forum for organizations with a shared purpose—CSIR.

On its “standards” web page, FIRST lists four standards it maintains:

- **Traffic Light Protocol (TLP)** on agreed-upon levels for marking information sensitivity
- **Common Vulnerability Scoring System (CVSS)** on describing the characteristics and severity of defects in software systems (not to be confused with **Common Weakness Scoring System (CWSS)** by the Mitre Corporation (MITRE))
- **Information Exchange Policy (IEP)** is a reporting format; in this regard it is another language for reporting, similar to IODEF or those listed in Table 6. Information Exchange Policy (IEP)’s focus is on disseminating information responsibly and quickly during CSIR.
- **passive Domain Name System (pDNS)** is a formatting standard for DNS traffic analysis; the FIRST group is working on an IETF standard.

None of these standards meet our relevance criteria, because none are about investigator practice. They are all things a competent investigator should know how to interact with and interpret, but they do not help us understand what decisions an investigator should make in a given scenario. FIRST also notes it contributes to several ISO standards, which Section 4.3.2 covers (namely, 27010, 27032, 27035, 27037, and 29147).

More instructive than these standards are FIRST’s “Security Reference Index” that is “helpful” to the FIRST community.[20] FIRST’s members are many, if not most, of the professionals and practitioners that we hope to come to understand. The documents Table 4 evaluates are listed as either best practices or standards in this reference index.[21] Five documents emerge as relevant to our review: Alberts et al. [2], Cichonski et al. [41], Fraser [61], Gorzelak et al. [66], Mundie et al. [132].

[20] https://first.org/resources/guides/reference.

[21] Strictly speaking, Mundie et al. [132] and Alberts et al. [2] are not linked directly; they are the most relevant part of a suite of publications linked to by FIRST as https://www.cert.org/incident-management/publications/index.cfm.

The education section of the FIRST website primarily contains one document, the services framework. There is a version for product security CSIR teams, but we will focus on the CSIRT services framework [16]. The services framework is a valuable document for defining what incident analysis is, and its relationship to incident management and other CSIRT services. However, the focus is on the relationship between a CSIRT and its constituency. Therefore, the descriptions of incident analysis are more like guidelines for service-level agreements rather than detailed instructions for how an incident analyst should perform the analysis or made decisions. Unfortunately, while the document does provide some useful broad constraints on the inputs to and outputs of analysis, the services framework does not directly bear on our research question of human decision-making during CSIR.

The relevant sections of the FIRST web page otherwise mainly link to the home pages of other security organizations; however, we cannot review the contents of everything these organizations have produced in full. In large part, the information is more about solving specific technical problems than our target for a general problem solving method. Such specific problems make for instructive cases when thinking about generalized methods, and so these organizations do provide an integral function to our topic. But they do not aim for the types of documents in scope of our review. The organizations identified are:

- **Center for Applied Internet Data Analysis (CAIDA)**, www.caida.org
- CERT/CC, www.cert.org
- **Center for Internet Security (CIS)** Benchmarking, http://www.cisecurity.org/
- Team Cymru, a security think tank, https://www.team-cymru.org/services.html
- **EU Agency for Network and Information Security (ENISA)**, https://www.enisa.europa.eu/ , including CSIRT services https://www.enisa.europa.eu/topics/csirt-cert-services
- **Open Web Application Security Project (OWASP)**, https://www.owasp.org/index.php/OWASP_Guide_Project
- Microsoft Security Guidance Center https://technet.microsoft.com/en-us/library/cc184906.aspx
- Sysadmin, Audit, Network, and Security Institute (SANS Institute) (SANS) reading room, https://www.sans.org/reading-room/

Although ENISA [54] targets management rather than practitioners, it provides links to training for practitioners. The two organizations the report lists are CERT/CC and the EU-funded TRANSITS. We handle CERT/CC above, but TRANSITS has not yet been mentioned.

TRANSITS is a training program administered by GÉANT with funding from ENISA. GÉANT is mainly responsible for managing the research network interconnecting 39 national research and education network institutions [64]. TRANSITS course materials are not public. Parties with a bona fide interest in conducting CSIR training courses can request permission for the introductory course, but not the advanced course [65]. Given these constraints, we have not been able to include the TRANSITS materials in this review. Similar access restrictions hold for CERT/CC course materials for "Foundations of Incident Management" and "Advanced Topics in Incident Handling"; similarly, these course materials are not included in our review.

Another important resource hosted by FIRST are the proceedings of FIRST events. The website hosts the agendas and descriptions of presentations as far back as 1996. At the 1996 conference, there were sessions on training analysts how to make decisions during incident handling [159]. Such tutorials remain common, for difference aspects of incident management, through to the 2019 conference. FIRST members have produced a wealth of case studies and tutorials about many topics, including incident management practices. The cases that are particularly effective or important might be expected to be represented in FIRST guidance such as Benetis et al. [16]. But there is no guarantee this is the case. Historically, FIRST guidance has been created in an ad hoc and regionally specific manner [158].

This situation of FIRST event proceedings poses a difficult methodology choice. On the one hand, any highly influential publications should be captured by our existing methodology, because they should be cited

Table 5. IC Results Summary

	Criteria				
Document	1	2	3	4	5
Heuer [79]	✓	✓	✓	✓	✓
Hutchins et al. [84]	✓	✓	✓	✓	✓
Caltagirone et al. [27]	✓	✓	✓	✓	✓
ATT&CK [169, 170]	✓	✓	×	✓	✓

The criteria are (1) target audience is security professionals; (2) topic in scope, per Section 2.1; (3) focus is investigator practices; (4) document finalized as of August 1, 2019; (5) available in English.

by in-scope standards. On the other hand, there is no guarantee of such referencing. There is not any easy assessment of which situation holds; there is no available, independent metric by which FIRST publications could be assessed for adoption or representativeness of practitioner behavior. For the purposes of this review, we take the hopeful hypothesis that standards will cite influential FIRST publications, and so our existing methodology should capture the salient ones. Future work should conduct a structured review of FIRST proceedings and publications to make this body of practitioner experience more accessible in general; however, we place that work out of scope here.

4.3.4 Intelligence Community. Table 5 summarizes the results. The canonical training course for CIA and other intelligence analysts is Heuer [79]. The book is essentially applied psychology. It covers topics such as analyzing competing hypotheses, which includes evaluating whether evidence has been planted to deceive, as well as overcoming human cognitive biases such as anchoring, vividness, and oversensitivity to consistency. Such methods, especially for evaluating evidence in the face of deception, have clear relevance to incident investigation.

The model of a computer attack as following a predictable "kill chain" of steps from start to finish was published by Lockheed Martin incident responders [84]. The seven steps are reconnaissance, weaponization, delivery, exploitation, installation, command-and-control, and actions on objectives. These are the steps in one single attack—a single phishing email, a single drive-by download with a malicious advert, and so on. Adversaries almost always compose a campaign out of multiple attacks; the objectives of one attack may be to obtain a platform from which further attacks are possible. The purpose of this model "is to capture something useful about the pattern all, or at least nearly all, attacks follow," so the analyst can anticipate what to look for or expect next [162, p. 10].

Caltagirone et al. [27] builds on attack ontologies, specifically the kill chain, and intelligence analysis to perform attribution in computer security incidents and analysis of whole campaigns. The method incorporates Bayesian statistics to model belief updates of the analyst. These statistical details are explicitly intended to help overcome analyst cognitive biases, such as those discussed in Heuer [79].

MITRE's **Adversarial Tactics, Techniques, and Common Knowledge (ATT&CK)** framework is work by a defense contractor that follows up on Hutchins et al. [84]. Adversarial Tactics, Techniques, and Common Knowledge (by MITRE) (ATT&CK) has a wider distribution and public profile than these other IC publications, but since it is a direct descendant of the kill chain it makes sense to discuss it here. ATT&CK is, first and foremost, an ontology. It organizes different ways that adversaries might achieve each of the steps in the kill chain [84]. However, it does not appear to make use of the diamond model, even though the realization of different kill chain steps changes based on where they are situated with the campaign. Either way, the use cases put forward for ATT&CK are adversary emulation, red teaming, behavioral analytics development, defensive gap assessment, SOC maturity assessment, and cyber threat intelligence enrichment [170, p. 3–4]. None of these are the practices

Table 6. Computer-security Related Reporting Formats and Data Formats Cited by IETF Standards Documents

Publisher	Type	Name
CERT/CC	Architecture	Automated Incident Reporting (AirCERT)
ICASI	Format	Common Vulnerability Reporting Framework (CVRF)
IEEE	Format	Malware Metadata Exchange Format (MMDEF)
IETF	Format	Intrusion Detection Message Exchange Format (RFC 4765) (IDMEF)
	Format	Incident Object Description Exchange Format Extensions (RFC 5901) (IODEF+)
	Format	RFC 5941, Sharing Transaction Fraud Data (extends IODEF)
ISO	Format	Software asset management: Software identification tag (ISO 19770)
FIRST	Data	Common Vulnerability Scoring System, maintained by FIRST (CVSS)
MITRE	Format	Common Attack Pattern Enumeration and Classification (by MITRE) (CAPEC)
	Format	Common Event Expression (by MITRE) (CEE)
	Data	Common Vulnerabilities and Exposures (by MITRE) (CVE)
	Data	Common Weakness Enumeration (by MITRE) (CWE)
	Data	Common Weakness Scoring System, maintained by MITRE (CWSS)
	Format	Malware Attribute Enumeration and Characterization (by MITRE) (MAEC)
	Format	Open Vulnerability and Assessment Language (by MITRE) (OVAL)
NIST	Data	Common Configuration Enumeration (by NIST) (CCE)
	Data	Common Configuration Scoring System (by NIST) (CCSS)
	Data	Common Platform Enumeration (by NIST) (CPE)
	Format	Open Checklist Interactive Language (by NIST) (OCIL)
	Format	Security Content Automation Protocol (by NIST) (SCAP)
	Format	Extensible Configuration Checklist Description Format (by NIST) (XCCDF)
XMPP	Format	XEP-0268 Incident Handling (using IODEF)

of the analyst or investigator during incident analysis. Organizing lower-level mechanisms certainly is helpful for understanding a higher-level mechanism (in this case, the kill chain) [163]. But ATT&CK, like other ontologies and reporting formats documented in Table 6, is not directly relevant to our topic of interest.

4.4 Referenced Documents

We harvest citations from the standards identified as relevant in Section 4.3. We summarize the results here; the evaluations are detailed in a subsection for each publication venue. The documents harvested from citations that are directly relevant to our review are:

- Carrier and Spafford [31]
- Casey [33, ch. 2]
- Ciardhuáin [40]
- Leigland and Krings [111]
- 27037 ISO/IEC [87]
- 27042 ISO/IEC [89]
- NIST SP 800-83 rev 1, §4 only [160]
- NIST SP 800-86 [102]
- Osorno et al. [137]
- Kossakowski et al. [104]
- Cheswick [35]
- Stoll [168]
- Mitropoulos et al. [127]
- Joint Chiefs of Staff [96, ch. 5 only]

4.4.1 Documents Referenced by IETF. The four relevant IETF standards reference 97 unique documents, excluding the IODEF-related standards already considered in Section 4.3.1. These documents fall into three broad categories: technical implementation requirements and dependencies; other related computer-security-incident report formats; and broader incident handling guidelines that describe the larger analysis and response context within which the reporting formats are used. This first category of implementation dependencies is not relevant to our project. Therefore, we focus on other reporting documents and broader incident handling guidelines. There are 22 cited reporting-related documents and exactly one related to broader incident handling and use of the reporting formats.

Table 6 lists the reporting formats and data sources cited by the IETF results in Section 4.3.1. Other report formats are primarily produced by NIST and MITRE—with funding from the US government including NIST. These projects also include the only referenced documents that are continuously updated data archives. The data formats for **Common Vulnerabilities and Exposures (CVE), Common Weakness Enumeration (CWE), CWSS, Common Configuration Enumeration (CCE), Common Configuration Scoring System (CCSS), Common Platform Enumeration (CPE)** are also continuously populated and published by NIST and MITRE as new vulnerabilities, platforms, and so on, are discovered or developed. Thus, these projects provide not only a format, but a standard reference dictionary of the possible terms with which the format may be populated. CVSS is perhaps the most important of these metrics that provide data and scoring; for a survey that relates Common Vulnerability Scoring System, maintained by FIRST (CVSS) to other security metrics, see Pendleton et al. [142].

These standard dictionaries are referenced by many of the other data formats that inherit the field essentially as a data type. For example, **Malware Attribute Enumeration and Characterization (MAEC)** may indicate which vulnerability a malware targets using its Common Vulnerabilities and Exposures (by MITRE) (CVE) number. Such dictionaries are useful background knowledge during CSIR, and help reduce confusion by providing common reference tags. However, following the pattern of other documents surveyed, these reference dictionaries do not provide agreed-upon evidence collection, field population, or analysis guidelines for their contents.

The next largest group of cited work from identified RFCs are three more IETF documents related to IODEF that did not appear in the original search. Other documents are also related to IODEF. CVRF and XEP-0268 extend and implement IODEF, respectively. AirCERT is a proposed implementation architecture that uses IODEF in automated indicator exchange. Our conclusions about the IODEF results identified in Section 4.3.1 therefore apply equally to these other documents, and we do not need to pay them special attention.

The remaining documents fall loosely into the NIST-MITRE orbit. ISO 19770 for asset management is developed separately from, but is related to, CPE and CCE, NIST's asset management for platforms and configurations, respectively. MMDEF is not directly related to MAEC; however, MAEC has adopted a significant component of the MMDEF schema.

The only citation related to actual decision-making during an incident is NIST SP 800-61 [41]. This document is already included from the FIRST results, see Section 4.3.3. Thus, from the IETF citations, we add no new documents to the review.

4.4.2 Documents Referenced by ISO. We extract the references from the three relevant ISO standards, namely, ISO/IEC 27035, ISO/IEC 27041:2015, ISO/IEC 27043:2015. ISO/IEC 27035 comes in two parts; although only the first part is in scope, we extract references from both parts. Although ISO charges for access to its documents, all the bibliographies are freely available, so we include all documents in this step.

There are 81 total references among the documents, with 64 unique references. Of these, 20 are elements of the ISO/IEC 27000 series of standards explicitly targeting information security. Four are the documents already referenced, and several others are already noted as not relevant in Table 3. However, from the references, we add the following two 27000-series publications to our survey documents as relevant to our survey (27037 and 27042):

- *Guidelines for identification, collection, acquisition, and preservation of digital evidence* [87]
- *Guidelines for the analysis and interpretation of digital evidence* [89]

Of the remaining 44 referenced documents, 13 are further ISO standards. Specifically:

ISO 15489-1
ISO 8601
ISO 9000
ISO/IEC 10118-2
ISO/IEC 12207:2008
ISO/IEC 17024:2012
ISO/IEC 17025:2005
ISO/IEC 17043:2010
ISO/IEC 20000
ISO/IEC 29147
ISO/IEC 30111
ISO/IEC 30121
ISO/IEC/IEEE 29148:2011

All of these other ISO documents are out of scope. We can further remove the following as unavailable or already evaluated: ILAC-G19, which directly follows from ISO 17020 and 17025; RFC 5070 [47], see Section 4.3.1; one by Valjarevic and Venter that is not available but appears by title and timing to be a working group presentation discussing the other two papers by these authors. We also will not consider the Daubert 1993 US Supreme Court case, as we aim to be jurisdiction neutral. Removing these leaves the documents listed and evaluated for relevance in Table 7. We pass the following documents on to the next stage of analysis: Carrier and Spafford [31], Casey [33], Ciardhuáin [40], Leigland and Krings [111]. And, like the IETF, ISO cites NIST SP 800-61 [41].

Also cited are two further MITRE data formats not covered in Section 4.4.1: **STIX (Structured Threat Information eXpression)** and **TAXII (Trusted Automated eXchange of Indicator Information)**. These build on MAEC, CVE, and so on, as formats for exchanging incident data. Like the other reporting formats already discussed, STIX and TAXII are not directly relevant to our CSIR decision-making topic. They are a language in which to do reporting; reporting is in scope. However, we plan to discuss what to report in a language-independent way.

The **ACPO (Association of Chief Police Officers)** guidelines are representative of many of the documents in Table 7. Their target audience is "UK law enforcement personnel who may deal with digital evidence" [187, p. 6]. It is primarily about the legal chain of custody necessary to bring digital evidence to court. This topic is about evidence collection, which is in scope. But the target audience is law enforcement, not security practitioners. The guidelines are not transferable to our topic of interest. The extent of comment on the actual work of understanding what the digital evidence means is constrained to "it is not practically possible to examine every item of digital data and clear tasking is needed to ensure that the digital forensic practitioner has the best chance of finding any evidence which is relevant to the investigation" [187, p. 10].

Some relevant work will not be carried through because it is obsoleted in a rather round-about way. Valjarevic and Venter [179, p. 1] notes "an effort to britishstandardise the process has started within ISO, by the authors." Thus, we consider papers by these authors to be obsolete because the authors directly subsumed their ideas into the ISO process. The process classes and activities used by ISO are clearly derived from Valjarevic and Venter [180, p. 6], which also contains a matrix of how these reference terms relate to other common forensic investigation ontologies. This set of works cited[22] matches the ISO work remarkably closely, as would be expected since the primary authors are the same. Unfortunately, Valjarevic and Venter [180] gives absolutely no methodology for how they arrived at this list of resources. Their analysis method is also not discussed, so it is unclear how or why they arrived at their categories and classification.

These omissions are particularly strange in that Valjarevic and Venter [180, p. 3] quote Cohen et al. [42] as rightly concluding the next steps in reaching consensus on and improving the field of digital forensics are a review of the literature that can be used to accurately drive consensus. This task is clearly what has been attempted, and as it has become an ISO standard it seems to have been accepted by a variety of practitioners. However, the lack of explanation of how these documents were selected as the correct set from which to drive consensus makes it hard to trace the authoritativeness of this source.

4.4.3 Documents Referenced by FIRST. As Table 4 demonstrates, we pass five FIRST-related documents through the evaluation of results to harvest further citations. Three of these documents do not have any

[22] Specifically, the overlapping works cited are Ballou et al. [11], Beebe and Clark [14], Carrier and Spafford [30], Casey [33], Ciardhuáin [40], Cohen [44], Leigland and Krings [111], Reith et al. [147], Williams [187].

Table 7. Documents Referenced by Relevant ISO Standards

	Criteria				
Document	1	2	3	4	5
Williams [187]	✗	✓	✓	✓	✓
Valjarevic and Venter [180]	✓	✓	✓	✗	✓
Valjarevic and Venter [179]	✓	✓	✓	✗	✓
Carrier and Spafford [30]	✓	✓	✓	✗	✓
Carrier and Spafford [31]	✓	✓	✓	✓	✓
Edwards et al. [50]	✗	✗	✗	✓	✓
Casey [33]	✓	✓	✓	✓	✓
Cohen [44]	✗	✗	✓	✓	✓
Cohen et al. [42]	✗	✗	✓	✓	✓
Palmer [138]	✗	✓	✗	✗	✓
Pollitt [144]	✗	✓	✗	✓	✓
Reith et al. [147]	✗	✓	✗	✓	✓
Beebe and Clark [14]	✗	✓	✓	✓	✓
Ciardhuáin [40]	✓	✓	✓	✓	✓
Leigland and Krings [111]	✓	✓	✓	✓	✓
Rowlingson [151]	✓	✗	✓	✓	✓
Hankins et al. [73]	✗	✗	✓	✓	✓
SWGDE [172]	✗	✗	✗	✓	✓
Garfinkel et al. [63]	✓	✗	✗	✓	✓
Ballou et al. [11]	✗	✗	✓	✓	✓
Alberts et al. [3]	✗	✗	✗	✓	✓
Cichonski et al. [41]	✓	✓	✓	✓	✓

The criteria are (1) target audience is security professionals; (2) topic in scope, per Section 2.1; (3) focus is investigator practices; (4) document finalized and not obsoleted as of August 1, 2017; (5) available in English.

citations ready to harvest. RFC 2196 [61] does not have in-line citations, and only Section 5.4 is relevant, so the relevant citations to follow cannot be distinguished. Further, RFC 2196 is already 20 years old, and so following any citations would provide little modern benefit. However, Gorzelak et al. [66] is a primary source—it is a survey of preventative practices at over 100 CSIRTs. Gorzelak et al. [66] notes the tools that the respondents use, but it makes no citations to other CSIR methodology documents. Alberts et al. [2] is similarly a primary source, though on incident management from CERT/CC. The only reference we take from Alberts et al. [2] is where it explicitly indicates further information on incident analysis is contained in another CERT document, namely, Kossakowski et al. [104]. Therefore, we harvest citations primarily from Cichonski et al. [41] and Mundie et al. [132].

Cichonski et al. [41] references three classes of resources. First is a list of CSIR organizations, second a list of NIST publications related to CSIR, and finally a list of applicable data formats. The list of organizations includes many already discussed in Section 4.3.3. Those jointly listed by NIST and FIRST are CERT/CC, ENISA, and FIRST itself. NIST additionally lists the **Anti-Phishing Working Group (APWG)**; **Computer Crime and Intellectual Property Section (CCIPS)**; **Government FIRST (GFIRST)**; **High Technology Crime Investigation Association (HTCIA)**; InfraGuard; the **Internet Storm Center (ISC)**; the National Council of **Information and Analysis Centers (ISACs)**; and **US Computer Emergency Readiness Team (US-CERT)**.

Table 8. NIST Publications Referenced by Cichonski et al. [41]

	Criteria				
Document	1	2	3	4	5
SP 800-53 [149]	✓	✗	✗	✓	✓
SP 800-83 (§4) [160]	✓	✓	✓	✓	✓
SP 800-84 [67]	✗	✗	✓	✓	✓
SP 800-86 [102]	✓	✓	✓	✓	✓
SP 800-92 [101]	✓	✓	✗	✓	✓
SP 800-94 [153]	✓	✗	✗	✓	✓
SP 800-115 [154]	✓	✓	✗	✓	✓
SP 800-128 [95]	✓	✗	✗	✓	✓

The criteria are (1) target audience is security professionals; (2) topic in scope, per Section 2.1; (3) focus is investigator practices; (4) document finalized and not obsoleted as of August 1, 2017; (5) available in English.

These organizations are certainly involved in various aspects of CSIR. However, organizations as such are out of the scope of our review.[23]

Table 8 lists and evaluates the relevance of the NIST publications referenced by Cichonski et al. [41]. All of these publications contribute to relevant background knowledge. For example, any CSIR professional will need to know what an intrusion detection and prevention system is and how they are deployed (SP 800-94). But this topic is not about evidence collection, analysis, and reporting; it is merely necessary background knowledge. The two publications that are relevant are the guides to *Malware Incident Prevention and Handling for Desktops and Laptops* [160, Section 4 only] and *Integrating Forensic Techniques into Incident Response* [102].

The data exchange formats listed by Cichonski et al. [41] are quite similar to those NIST, IODEF, and MITRE formats extracted from the IETF documents in Table 6. The only difference is the addition of **Asset Identification (AI), Asset Results Format (ARF)**, CVSS (from FIRST, see Section 4.3.3), and **Cyber Observable eXpression (CybOX)**. As discussed in Section 4.4.1, these formats are languages for reporting results, but they do not directly discuss what to say. Our intention is to discuss what to say in results, while being language agnostic, which puts these various formats and languages just out of scope.

Mundie et al. [132] cites 27 documents. They include several technical format documents for ontologies in the W3C **Ontology Web Language (OWL)**, KL-ONE knowledge representation, knowledge graphs, **process specification language (PSL)**, or the display tools used (Graphviz), which we will not discuss. There are also psychology and ontology that are obviously out of scope for our review [9, 124]. Further, four references we have already considered, namely, ISO/IEC 27001, ISO/IEC 27002, Cichonski et al. [41], and Beebe and Clark [14]. These exceptions leave the eight documents evaluated in Table 9. The only documents that pass the evaluation are Osorno et al. [137] and Kossakowski et al. [104].

MITRE Corporation [126] on whether cybersecurity is a science is not within our current, relatively narrow scope. However, since CSIR is an important subset of cybersecurity, whether security investigations are a kind

[23] As a convenient sample, one of the authors has presented at Anti-Phishing Working Group (APWG) [161, 164] and attended InfraGuard meetings, and does not expect there would be significant benefit in expanding the scope to include them. Likewise, the same author has interacted with several Information Sharing and Analysis Centers (ISACs), and reviewed their available materials (Research and Education Networking ISAC (REN-ISAC) and Financial Services ISAC (FS-ISAC) especially) and does not believe they have any documents of importance to our topic.

Table 9. Documents Referenced by the CERT Documents

	Criteria				
Document	1	2	3	4	5
MITRE Corporation [126]	✗	✓	✗	✓	✓
Mundie and Ruefle [131]	✗	✓	✓	✗	✓
Fenz and Ekelhart [57]	✗	✓	✗	✓	✓
Osorno et al. [137]	✓	✓	✓	✓	✓
Magklaras and Furnell [118]	✓	✗	✓	✓	✓
Wang and Guo [185]	✓	✗	✗	✓	✓
Chiang et al. [37]	✓	✗	✗	✓	✓
Ekelhart et al. [53]	✓	✗	✓	✓	✓
Kossakowski et al. [104]	✓	✓	✓	✓	✓

The criteria are (1) target audience is security professionals; (2) topic in scope, per Section 2.1; (3) focus is investigator practices; (4) document finalized and not obsoleted as of August 1, 2017; (5) available in English.

of subcategory of scientific investigation clearly impacts our handling of what CSIR is and how to link it to knowledge generation and evidence evaluation more generally. Spring et al. [165] address the relationship between cybersecurity and science and argue that cybersecurity as practiced is a kind of science.

Fenz and Ekelhart [57] provides a difficult decision. It is one of the few attempts at formalization. However, its target is security knowledge, not security practice. This topic is closely allied to our hope to formalize CSIR, as security knowledge would be instrumental to that project. So while not in scope for this review, this document may be useful for future related work.

4.4.4 Documents Referenced by the IC. Heuer [79] presents some challenges to adequate reference harvesting. The book contains no collected list of references, written in a traditional humanities style in which references are in footnotes intermixed with commentary, but this is not the central problem. As essentially a military intelligence and psychology book, its sources are quite wide-ranging. References range from World War II Nazi-propaganda analysis to behavioral economics. It is only through Heuer's CIA experience that these disparate sources are converted into a useful guide on how to reason in adversarial situations. The other challenge is that Heuer [79] makes only passing reference to computers as tabulating machines. The closest he seems to get to computer science is via Simon [157], as he discusses decision-making and satisficing. For these three reasons, we consider Heuer [79] as essentially a primary source and do not trace citations from it. One should not be surprised it has many features of a primary source, as surely its main value is summarizing CIA analytic experience not otherwise publicly available.

Caltagirone et al. [27] and Hutchins et al. [84] are more straightforward. There are 79 citations between the two, with no overlap, though Caltagirone et al. [27] cites both Hutchins et al. [84] and Heuer [79]. The references in Caltagirone et al. [27] are noticeably more strategy-focused over the tactically focused Hutchins et al. [84], as one would expect from their different topics. Hutchins et al. [84] cites several vulnerability bulletins and company advisories as cases; it is more of a primary source, documenting the analysis methods used by Lockheed Martin CSIR staff. We do not consider such advisories, software tool documentation, and news items, as they are not within our review topic. There are also several references already covered elsewhere: Cichonski et al. [41], STIX, CVE, and SANS. There is even yet a new reporting and data exchange format: **Vocabulary for Event Recording and Incident Sharing (VERIS)**. We also exclude references that are merely the official definitions of terminology. These exceptions reduce the total referenced works to 50.

Table 10. Documents Referenced by Caltagirone et al. [27] and Hutchins et al. [84]

	Criteria				
Document	1	2	3	4	5
Stamos [167]	✓	✗	✗	✓	✓
Amann et al. [5]	✓	✗	✗	✓	✓
Cheswick [35]	✓	✓	✓	✓	✓
Stoll [168]	✓	✓	✓	✓	✓
Duran et al. [49]	✓	✗	✗	✓	✓
Cohen [43]	✓	✗	✗	✓	✓
Lewis [112]	✗	✗	✗	✓	✓
Tirpak [175]	✗	✓	✓	✓	✓
Hayes [76]	✓	✗	✗	✓	✓
Willison and Siponen [188]	✓	✓	✗	✓	✓
Mitropoulos et al. [127]	✓	✓	✓	✓	✓
Caltagirone and Frincke [26]	✓	✗	✓	✓	✓
Caltagirone [25]	✓	✗	✓	✓	✓
Bellovin [15]	✓	✓	✗	✓	✓
McClure et al. [120]	✓	✗	✓	✓	✓
Brenner [21]	✗	✗	✗	✓	✓
Van Eck [181]	✓	✗	✗	✓	✓
John and Olovsson [94]	✓	✗	✓	✓	✓
Joint Chiefs of Staff [97]	✓	✗	✗	*	✓
Joint Chiefs of Staff [98]	✓	✗	✓	*	✓
Joint Chiefs of Staff [96, ch. 5 only]	✓	✓	✓	*	✓

The criteria are (1) target audience is security professionals; (2) topic in scope, per Section 2.1; (3) focus is investigator practices; (4) document finalized and not obsoleted as of August 1, 2017; (5) available in English. The (*) in criterion four indicates we are referencing an updated version of the document.

The kill chain [84] and the diamond model [27] are both attack ontologies. They model the possible routes an adversary may take when executing an attack. One big class of cited work is other attack ontologies. We consider the kill chain and diamond model as *de facto* standard ontologies, but we believe they reached that level of agreement within the IC, because they also come with investigative norms for interpreting and filling in the ontologies. There are many other attack ontology works that do not come with such guidance, and so we will not consider them directly in scope for our review. We remove 16 references from consideration, because they are attack-ontologies without guidance.

There are also several references that are clearly for background or motivation, such as Hawkins [75], Liu and Motoda [115], Symantec's analysis of the Duqu malware, and assessments of Chinese attack capabilities. We also remove how-to descriptions for conducting particular methods of technical analysis, namely, on passive DNS analysis [7], crime-pattern analysis [139], and honeypots via the Honeynet Project. This leaves us with 22 documents in Table 10. Of these, four pass our relevance requirements: Cheswick [35],[24] Stoll [168], Mitropoulos et al. [127], and Joint Chiefs of Staff [96, ch. 5 only].

[24] Technically the citation is to this article's republication in a popular textbook, Cheswick et al. [36, ch. 16]. However, as the rest of the textbook is not directly referenced, we reference just the original publication.

Table 11. Categorization of Relevant Documents

	Directness per Phase			Scope per Goal					
Document	Col	Anz	Rep	Fix	Int	LE	General	Type	Formal
27035-1:2016 [91]	C	C	C	B	×	×	Un	Ont	Qual
27037 ISO/IEC [87]	D	×	×	×	×	M	Un	Ont	Qual
27041 [88]	C	C	C	×	×	×	Likely	Instr	∅
27042 ISO/IEC [89]	×	D	D	×	×	M	Un	Instr	∅
27043[90]	C	C	C	B	×	M	Un	Ont	Qual
RFC 2196, §5.4 only [61]	C	×	D	M	×	B	Likely	Instr	∅
RFC 6545 [129]	C	C	D	N	B	×	Un	Ont	Formal
RFC 7203 [173]	C	×	C	N	N	N	Un	Ont	Formal
RFC 7970 [46]	C	×	D	M	M	M	Un	Ont	Formal
RFC 8134 [85]	C	×	×	B	B	×	Un	Study	∅
NIST 800-61 [41]	C	D	C	M	×	M	Un	Adv	Qual
NIST 800-83 §4 [160]	D	×	C	B	×	×	Un	Ont	Qual
NIST 800-86 [102]	C	×	C	×	×	×	High	Study	Qual
Gorzelak et al. [66] (ENISA)	D	×	×	×	N	×	Un	Study	Qual
Alberts et al. [2]	×	×	C	B	×	×	Un	Ont	Qual
Kossakowski et al. [104]	C	D	C	M	×	×	Likely	Adv	Qual
Mundie et al. [132]	C	C	C	B	×	×	High	Ont	Perf
Osorno et al. [137]	×	C	C	B	B	×	High	Ont	Qual
Hutchins et al. [84] (IC)	C	C	×	×	M	×	High	Ont	Qual
Caltagirone et al. [27] (IC)	C	D	×	×	M	×	High	Adv	Perf
Heuer [79] (CIA)	×	D	×	×	B	×	Wide	Instr	Qual
JCS [96, ch. 5 only]	×	D	×	×	M	×	Likely	Instr	Qual
Casey [33, ch. 2]	C	D	D	×	×	B	Wide	Ont	Qual
Mitropoulos et al. [127]	C	C	C	M	N	×	Un	Study	Qual
Carrier and Spafford [31]	D	C	D	×	×	M	Likely	Ont	Qual
Ciardhuáin [40]	C	C	C	×	×	B	Un	Ont	Perf
Leigland and Krings [111]	D	×	×	N	N	N	Likely	Instr	Formal
Luttgens et al. [117]	D	C	D	B	×	N	High	Adv	∅
Stoll [168]	C	D	C	N	M	N	Likely	Study	∅
Cheswick [35]	C	D	×	×	B	×	Likely	Study	∅

The phases are collection, analysis, and reporting. Cells have a cross if a phase is not addressed. Advice directness values are direct (D) or constraints-based (C). Values for goals are to fix an infected system (fix), gathering intelligence (int), and law enforcement (LE) action. Values for a document's intended scope are narrow (N), medium (M), or broad (B). Values for generalizability of an approach are unlikely (Un), likely (Likely), highly likely (High), or already widely generalizable (Wide). Document types are case studies (Study), ontologies (Ont), advice on actions (Adv), and explicit instructions (Instr). Values for formalization are not present (∅), qualitative (Qual), formal, or perfunctory (Perf). For more information about the construction of this table, see Appendix A.

5 SYNTHESIS

The search, appraisal, and reference harvesting stages of this review have returned 29 documents for further analysis from roughly 350 possible documents returned from search. Section 5.1 presents the values and concepts by which we organize the relevant documents. Section 5.2 presents the three gaps we identify in the CSIR documents. Sections 5.2.1 through 5.2.3 use the concepts from Section 5.1 to identify which documents would reasonably be expected to fill each respective gap, and then discusses each document in turn to show how it does not fill the gap.

5.1 Document Classification Methodology

Section 4 has collected the relevant documents about investigator practice during evidence collection, analysis, and reporting. Our synthesis goal is to evaluate the nature and quality of advice that these documents provide about making decisions during these phases of CSIR.

To assist our synthesis and discussion, we classify advice on these topics in the ways listed below. Using this classification will allow us to assess: to which phases the document applies, the directness with which the document applies to each phase, the type, investigative goals supported, broadness of scope, generalizability of advice, and formalism. We use this initial evaluation to identify groupings of documents and get an overview of what the literature search has found to be available. We group the documents by gaps in the literature, but use the phases and type particularly to group documents around those gaps for discussion.

Phases indicates simply what combination of evidence *collection*, *analysis*, and *reporting* the document covers.

Directness has two possible values: *direct* and *constraints*. Direct commentary on CSIR explicitly talks about what an investigator should or should not do. Constraints provide only requirements for outcomes or outputs and do not indicate how these properties should be achieved. Constraint-based advice is common when situating investigation within the larger context of CSIR, and situating response within incident management.

Type indicates what type and level of detail the document provides to decision-making. Possible values are *case study*, *ontology*, *advice*, and *instructions*. At one end of the spectrum are case studies. Case studies report the facts of an individual case of investigation, without attempting to abstract up to general lessons. A categorization forms categories of useful actions (implicitly or explicitly from case studies), but gives no advice on how to apply these ontologial categories. Advice provides some ordering on what category of action should be taken, given certain conditions. Finally, instructions provide explicit decision-making instructions on how to evaluate options. Type also provides some rough guide to how much effort it will take to apply the document to practice, with case studies being the most difficult.

Goals indicates what sort of investigation the advice targets. We distinguish three goals an investigator could have: *fix* the system, gather *intelligence* on adversaries, and make *arrests*. Certainly, there may be other goals, but these cover a wide degree of practical differences. Investigators need quite different information between these goals. For example, to fix a system, one needs to know everything that has been accessed by the adversary, but you need to know rather little about them. Whereas to make arrests, one cares very much about the adversary, but also is bound by several practical matters of what counts as admissible legal evidence of attribution and loss. When gathering intelligence on what an adversary may do next, these legal considerations fall away, but one also focuses on quite different aspects than fixing a system. For example, to gather adequate intelligence one need not enumerate all compromised systems.

Scope reports how widely the document applies, as reported by the document. Options are *narrow*, *medium*, and *broad*. A narrowly scoped document is intended to apply to only a small, non-representative group of people and/or for a short period of time. Broad scopes are intended for most people within information security. Medium scope fits somewhere in between. Examples of medium scope are US-based law enforcement forensics specialists, or the operators of tier-three (that is, backbone) networks.

Generalizability of advice indicates how likely it is that the document can be relevant to contexts outside those for which it was specifically designed. Generalizability is explicitly level-set from the document's scope. Thus, a document with broad scope but no generalizability may still be applicable to more people than a narrowly scoped document that is generalizable. Whereas scope is a measure taken directly from the document being evaluated, generalizability is an evaluation of potential not explicit in the document. Indicators of generalizability include use of models or methods from other disciplines with well-established other uses or evidence from sources other than the document itself that the advice from the document applies more widely. Options for this criterion are coarsely set as *unlikely* (<15%, ±5%), *likely* (in between unlikely and highly likely), and *highly likely* (>85%, ±5%); values represent essentially the evaluator's prior belief on the document being applicable outside its stated scope. We have a final value, *widely*, indicating the document is certainly generalizable beyond its scope and is likely to be generalizable to a much broader scope.

Formalism reports the degree of mathematical formalism present in the advice provided. Options are *none*, *qualitative*, *formal*, and *perfunctory*. Perfunctory indicates formalization is present, but essentially unused to advance the arguments or positions of the document. This rating does not mean the formalism is wrong; however, it does indicate it would take significant effort on the part of the reader to make use of the formalism beyond what qualitative models would provide. None only applies to narratives that make no attempts at abstraction. Both qualitative models and formal mathematical models have value in their own ways, and one should not be considered preferred over the other per se.

Section 5.2 synthesizes the results after this classification of the documents. Our focus is on how investigators make decisions about evaluating the quality and importance of evidence, generalize from particular evidence to evidence of trends or patterns of behavior, and select what to report based on security constraints as well as what others will find most convincing. Although these align loosely with the three phases of investigation we have highlighted, we are not making a one-to-one connection between the three phases and our three focal points. For example, if an investigator knows some kind of evidence is particularly convincing to report, that should impact what they look for during the evidence collection phase.

5.2 Synthesis Results

Table 11 classifies advice on the 30[25] selected documents in several ways: to which phases the document applies, the directness with which the document applies to each phase, the applicability of the advice, investigative goals supported, broadness of scope, generalizability of advice, and formalism. We shall make some commentary on all the documents in the table, in order of appearance, with one exception. Before we lose ourself amongst the trees of this discussion, we will make our general claim of the primary gap in the literature clear up front.

Kent et al. [102, p 3–8] recommends all organizations have an CSIR capability and that analysts use "a methodical approach to a digital forensic investigation." Despite recommending everyone use a methodical approach, NIST fails to provide one. This failure is symptomatic of the state of available practicable policy advice and practitioner training material. This is the central gap identified by the literature review: There may be adequate concrete tools and training available, but there is no general training or policy advice for strategic selection of tactics, that is, *which analysis heuristic or technical tool to employ in a particular situation and why*. An attendant gap is a failure to advise on *when the analyst is justified in generalizing*; that is, making a stronger, broader claim from singular pieces of evidence. Because there is no advice on which strategy to employ, or when broadening claims are justified, there is similarly a gap in *what information to report to convince a reader that the analyst should be believed*.

Kent et al. [102] clearly states the importance of digital forensic investigation and advises on terminology, analysis techniques, and pitfalls to avoid. Of any document in Table 11, NIST SP 800-61 comes closest to providing a methodical approach, but the analysis method there still amounts to an unordered collection of tips, tricks, and pitfalls to avoid. The NIST documents evaluated and surveyed here are a generally positive attempt at providing practical advice to a wide audience on a complex topic. But the assumption is that once an investigator is taught how to use a tool, they will know when and why to use it. This gap is a recurring assumption, and precisely our intended focus for improvement.

5.2.1 Selecting an Analysis Heuristic or Technical Tool. None of the documents in scope provide advice on which analysis heuristic or technical tool to employ in a particular situation and why. That is, there are no

[25] As a result of DTRAP reviews, we added Luttgens et al. [117]; see Appendix C.

hypothesis generation guidelines for CSIR. Table 11 contains some documents that present either advice or instructions (as the type) directly (D) about the analysis phase, namely, References [27, 41, 79, 89, 96, 104]. We also might expect to find this kind of advice from documents that provide advice or instructions about the collection of evidence; only Leigland and Krings [111] fits those criteria. Unfortunately, none of these documents contain this level of detail about how to manage analyst resources and how to direct analyst inquiries. Each document presents some useful advice or instructions about how to do analysis, but the scope is either too broad or too narrow. We discuss what each document provides in turn, in the order it appears in Table 11, to sketch the absence of advice on selecting an analysis path.

ISO/IEC 27042 [89] provides a basic distinction between static and dynamic analysis of malware (it uses "live" for dynamic), but all that is really provided are a few descriptions of what distinguishes static and dynamic analysis. These descriptions do not provide information on how to actually do either kind of analysis, or even common pitfalls or errors to avoid.

NIST SP 800-61 [41] is cited by all four venues, and some use it as their standard directly, with good reason. It is comprehensive and thorough without being overbearing. Indeed, Reference [41] comes closest to providing guidance on analysis heuristics, in that it recognizes the problem. However, its focus is incident management, not investigation. The analysis phase receives about three pages of discussion (28–30), reporting one page (31),[26] evidence collection half a page (36), and general decision-making and prioritization two pages (32–33). NIST SP 800-61 [41, p. 32] addresses the problem of scarce resources directly: "prioritizing the handling of the incident is perhaps the most critical decision point in the incident handling process." The following discussion, while short, is two more pages about decision-making during incident analysis than almost any other document found. Regardless, it is not sufficient to develop a robust account of the nuances and difficulties an investigator regularly deals with when evaluating evidence, generalizing from particulars, and deciding how best to report their findings.

Kossakowski et al. [104, p. 17ff.] provides classes of advice, like collect logs and isolate infected machines from the network. These perhaps come the closest to advice about how to collect evidence from computer incidents. However, it is silent on which logs to collect, or what to look for when examining network traces. While this advice is highly likely to be able to generalize to all cases of CSIR, the level of detail is not operationalizable as decision-making instructions.

Caltagirone et al. [27] provides some light formalization of their qualitative categories of the structure of an adversary campaign into both graph theory and subjective probabilities and Bayesian statistics. The structures are constructed adequately; however, any application of them is left as an exercise for the reader. The main achievement of the document is an analysis structure that is a way to process information once the analyst has it. Caltagirone et al. [27] helps an analyst make sense of what information they already have. This assistance is certainly useful, but it does not help an analyst decide what they need to know, how to seek it out, or how to prioritize what to look for.

Heuer [79] is an instructive case for some advice being too broad for our purposes. The book is comprised of explicit decision-making instructions for analysis of intelligence. However, the level of abstraction is so broad that it applies to almost any adversarial decision-making environment. The advice is valuable, especially the advice for avoiding cognitive biases, and any advice for analysis within CSIR should be consistent with Heuer [79]. However, it does not provide instructions at the level of detail that is directly useful for CSIR.

Joint Chiefs of Staff [96, ch. 5] is about how to think like your adversary. It is an extended treatment of developing and evaluating adversary action plans across multiple dimensions under constrained resources.

[26] SP 800-61 acknowledges its discussion of reporting is too brief and refers the reader to RFC 5070. This document has since been obsoleted by RFC 7970 [46].

The basic cycle is to identify objectives, enumerate courses of action, evaluate and rank the likelihood of following each action, and identify necessary intelligence collection requirements to determine adversary decisions. The document is about military intelligence operations generally, not computer-security incidents. However, one narrow but necessary aspect of any investigation is how to anticipate an adversary. This document covers the thought process behind the topic of anticipation in a way that should be applicable to computer security. Similar to our analysis of Heuer [79], any advice for analysis within CSIR should be consistent with Joint Chiefs of Staff [96, ch. 5]. But it does not contain positive instructions about how an analyst might decide what tool or thought process to employ in a particular CSIR situation.

Leigland and Krings [111] provides a formal specification of evidence collection methods that can be adapted to specific operating systems. The language associates collection goals with certain common attack patterns. The goal is narrowly practical—to speed collection of evidence by technicians during an investigation while reducing superfluous data collection to make analysis a bit easier. The language maps general actions to specific operating-system commands. The technician needs to specify file identifiers for specific attack campaigns; thus, methodology requires an up-to-date and accurate blacklist of file identifiers. As with blacklists for network touchpoints, we suspect this approach is unsustainable [122]. Furthermore, the approach merely obfuscates or defers our question of interest—one cannot create a good blacklist without past incident analyses. Therefore, this document is more accurately about how to efficiently use past analysis results, rather than how to prioritize analyst choices in a present CSIR.

Based on our detailed analysis of these documents, there are no CSIR standards that provide advice on which analysis heuristic or tool to use at one time or in what situation, given limited analyst resources. The search results and categorization of documents captured in Table 11 suggests that the above documents are all the relevant documents to consider. Therefore, while we have found useful related documents, we are confident this is a bona fide gap in CSIR standards.

5.2.2 When Generalizing is Justified. By "generalizing," we mean making a stronger and/or broader claim from singular pieces of evidence. Breadth takes many forms, such as concluding that claims hold across longer periods of time, more types of devices, higher number of victims, or connecting previous disparate events to a common source. Generalizations do not have to be grand to be of great importance to CSIR. If an analyst observes evidence of compromise on one end-user host, then their conclusions are often limited. But deciding when the analyst should generalize from "just the observed compromises occured" to "the whole network is compromised" (including hosts for which direct evidence is not available) is important. Another example is when blocking network traffic. Sometimes an analyst sees enough malicious traffic from a set of 256 IP addresses to block all of them, even those that have not yet sent malicious traffic. We are seeking advice on deciding what counts as "enough" in this situation, as well as why to choose 256 addresses (a /24 netblock) rather than 128 or 512. There is no simple rule, so advice would be important.

From the categories presented in Table 11, any document that bears on this problem would need to have the analysis phase in scope. Of these, case studies may contain the space and detail to at least demonstrate how past analysts have handled this issue. The relevant case studies are References [35, 127, 168]. An ontology document may propose or maintain relationships between concepts that guide generalization. The ontology documents with an analysis-phase scope as their most direct scope are References [33, 40, 84, 90, 91, 132, 137]; we will check these documents for guidance about when an analyst is justified in generalizing. The documents that offer advice or instructions directly on analysis, discussed in Section 5.2.1, also need to be examined as potentially relevant.

Table 11 also reports our assessment of the generalizability of the document itself. This assessment is not about whether the document contains advice about how to generalize, but rather is about how broadly applicable we find the document to be. That column is unrelated to our present discussion.

We discuss the candidate documents in the order they appear in Table 11.

ISO/IEC 27035-1 [91] and 27043 [90] are not useful structures for organizing knowledge, let alone how or when to generalize. As discussed further in Appendix A, the documents contain jarring transitions of terminology and concepts that frustrates any attempt to use them as a genuine organization of knowledge about CSIR.

Mundie et al. [132] is, in effect, a literature review of incident management. As such, it mostly constrains the inputs and outputs one would expect from CSIR. The formalism provided is in a specification of an OWL ontology language of incident management. This language is a useful step in reconciling various incident management processes. However, it is a much more abstract than our current question.

Osorno et al. [137] has done something similar to our project here, in that they inventory various incident management processes, with two main differences. They focus on moving up a level to inter-organizational coordination during complicated incidents, rather than zooming in on individual analyst decision-making. Osorno et al. [137] also focuses on the US context. This different purpose leads to substantial differences in emphasis as to what is reviewed; for example, where we have generally set aside data exchange formats (see Section 4.4.1), Osorno et al. [137] spend considerable effort mapping these formats into each other. For this reason, the extent of their recommendation on CSIR amounts to do an OODA-style loop, a military term standing for observe, orient, decide, and act [137, p. 7]. This advice is not incorrect, but the question of how to orient and how to decide is precisely the gap we are identifying, and this document does nothing to fill it in.

Hutchins et al. [84] presents an ontology of adversary action sequence, based on eight years of case studies. The authors use this ontology to advise network defense. However, network defense is not CSIR. The level of advice is on the order of "to disrupt installation events, use anti-virus." This advice is sound, but it will not help a CSIR analyst generalize. Spring and Illari [163] argues that the kill chain model from [84] could be situated as part of a set of generalization heuristics, but [84] itself does not discuss how the kill chain model might be used to make decisions during CSIR.

Casey [33, ch. 2] is built around the claim that digital forensics is just another kind of scientific investigation. The target audience is law enforcement who will be using information technology to support general legal cases, not computer crimes. The scientific method is represented as simply create and evaluate hypotheses dispassionately based on evidence. This description is supported by several case studies as examples working through the method. The shortcoming here is not in Casey [33] per se, but in the naïve scientific method put forth. One major shortcoming is that there is no discussion of how to create hypotheses nor any process for deciding when the resulting evidence generalizes beyond the context of the observed evidence. Evidence in cybersecurity is complex and context-dependent [163], but more general guidance should be offered in terms of methods for finding key pieces of evidence that work on groups of related cases. This would extend more widely than individual case studies, but be much less abstract than "evaluate hypotheses using evidence" as suggested by Casey.

Luttgens et al. [117], especially Chapter 11, continues in the tradition established by Casey [33] that digital forensics (and therefore CSIR) is a kind of scientific investigation. They provide some very broad categories of analysis methods and some technical examples. The methods they introduce briefly are "use of external resources, manual inspection, use of specialized tools, data minimization through sorting and filtering, statistical analysis, keyword searching, [and] file and record carving." The authors combine this advice with case studies that make the advice more useful than this simple list appears. However, the authors appear to put a good deal of weight on assuming their readers agree on what being "scientific" means. This assumption does not appear to be warranted within the cybersecurity community—there is active disagreement on that term [165]. Luttgens et al. [117], like Casey [33], is an excellent starting place for CSIR professionals. But neither teach a practitioner why one method is chosen over another, how to prioritize amongst methods given limited time, or systematize their decision-making.

Mitropoulos et al. [127] aims at providing an account of CSIR with adequate detail for evaluating whether an incident management team within an organization is functioning properly. To achieve this, the authors specify several examples of how an incident should be handled. For example, the authors provide a flow chart of analysis of IDS logs, intended to serve as an example of analysis. This result appears promising, it is directly a representation of reasoning and decision-making during CSIR. However, because the target audience is management of CSIR functions, Mitropoulos et al. [127] falls short of advising an analyst on how to answer the questions or generalize results. The authors note that these things should happen, and where in the CSIR process they should happen. But mostly the authors enumerate the available options, rather than guide how to select among them.

Ciardhuáin [40] attempts a novel formalization by incorporating information flow; although not cited, this is a term likely taken from Barwise and Seligman [13]. Information flow is a reasonable choice for how one might model generalization. However, the application to forensic investigation [40, p. 21] bears little resemblance to formal information flow models. The discussion places vague constraints on CSIR, such as "the investigators must construct a hypothesis of what occurred" and "[t]he hypothesis must be presented to persons other than the investigators" [40, p. 7]. But in practice these constraints are of little use.

Stoll [168] and Cheswick [35] are two of the earliest detailed case studies of a CSIR analyst, from a first-person perspective of what the analyst was thinking and why they took certain actions. Modern case studies often do not discuss how exactly the analyst found what they found. These omissions are with good reason; adversaries are likely to read any reports their targets publish. Paradoxically, this makes the old case studies more valuable, as they remain some of the better expressions of the analyst's thought process. The fact that Stoll [168] and Cheswick [35] are cited by standards 20 or more years after the case studies were published attests to this value. The tools and networks the old case studies discuss are outdated, which can make them hard to apply to modern systems. Despite this, the insights into analyst processes are probably the closest thing to what we are seeking insofar as advice on when an analyst is justified in generalizing. One difficulty is that, while Stoll [168] and Cheswick [35] each justify the generalizations they make, they do not provide much guidance on how to do something similar in other situations. Stoll [168] explicitly cites his training as a physicist and calls his approach scientific; this is better than Cheswick [35], which provides no explicit discussion of this topic. However, the description of this method in Stoll [168] is naïve in ways similar to that of Casey [33]. And so for similar reasons as explained above in relation to Casey [33], these case studies do not quite provide the guidance CSIR analysts should have.

Section 5.2.1 revisited. The documents that fail to provide advice on what analysis tool when also tend to fail to provide advice on when generalizing is justified. Several of them might be useful, but none of them address generalizing directly. The diamond model has the same problems as Hutchins et al. [84], discussed above. Heuer [79] is the best available advice, in that it discusses situations in which generalization is not justified due to cognitive biases interfering with the generalization reasoning. But Heuer [79] does not distinguish between generalization and other forms of reasoning, and because his target audience is intelligence analysis generally, challenges in CSIR that are specific to computer science or networking are not addressed. The other papers Section 5.2.1 discussed basically do not address the analysis method, and, since generalization strategies would be part of an analysis method, this gap remains.

5.2.3 Convincing Reporting. This gap is a lack of advice on what information to report to convince a receiver that the analyst should be believed. Our analysis separates the receiver believing a report and taking action based on the report, but in the pragmatic world of CSIR, a report should enable the receiver to take action based on this new belief or information. From Table 11, we will discuss those that directly bear on the reporting phase. The ISO document, 27042, suffers from the same general poor quality as the other 27000-series documents (see Appendix A), and we will not discuss it in detail. The IETF documents are focused around RFC 7970, so we will only discuss that one in detail. RFC 2196 is not related to it, but it is 20 years old and it has little modern use.

Finally, Casey [33] is related, insofar as it relates this kind of reporting to scientific reporting. As discussed in Section 5.2.2, the failing is a naïve sense of what scientific reporting means. We fully endorse the analogy to scientific reasoning, but to make it work, CSIR needs a better integration with what scientific practice actually is, rather than a cartoon version of it. This leaves two documents about reporting to discuss in detail to explain why this gap remains in the literature: Danyliw [46] and Carrier and Spafford [31].

RFC 7970 [46] is the heart of the IETF incident analysis standardization effort. While a useful structure for reports, it does not discuss why to report certain things over others. However, it is widely influential; this RFC obsoletes RFC 5070, which is cited or used by most publication venues as the incident reporting format. The focus is on exchanging indicators of incidents for collective defense. Although IODEF is, strictly speaking, just an XML schema for document incidents, the available options and the ontology provided to constrain the other phases up to reporting. For some fields, this provides only minimal collection requirements. However, consider the system impact type attribute, which is a required field. There are 24 options specified, ranging from "takeover-account" to "integrity-hardware" [46, p. 46]. Individuating among these various impacts would require a relatively sophisticated CSIR and analysis capability; it is not so easy as logging an IP address and passing it along. Just within the assessment class, one of two dozen overarching classes, there are five types of impact to distinguish between with similar detail: system, business, time, money, and confidence. Such detail provides the most rigorous reporting requirements and guidance available. The other RFCs cited are all extensions of IODEF, which are useful to varying degrees, but do not qualitatively change our analysis. IODEF constrains CSIR reporting in useful ways, but it does not advise on importance or utility of what information items a report is constrained to.

Carrier and Spafford [31] describes an evidence collection and hypothesis testing and reporting model of digital forensic investigation. The insight that CSIR and digital forensics is more like crime scene investigation than traditional police forensics is invaluable and helps deconflict jargon between fields of study. Their target goal is gathering of adequate legal evidence, and so this would seem quite promising as far as addressing what to report, at least in some subset of goals in CSIR. However, their advice amounts merely to saying that deciding on collection and reporting targets "is the most challenging [aspect] of the search phase" and is done "from either experience or existing evidence" [31, p. 8]. Like the other documents, this one sidesteps most of the gaps we identify as essentially out of their scope. It is possible that the metaphor to crime scene investigation was meant to open up a different literature that would fill gaps such as this; however, those connections are not explicit and so we are left with a gap.

These three gaps leave a CSIR analyst in a difficult position. Without advice on what evidence to collect, how to synthesize it, and what to report, they are left to analyze essentially as a trial by fire. This unguided analyst training is unlikely to be successfully scale-able, even though there are certainly some analysts who succeed and do the job well.

6 CONCLUSION

Our review of the incident analysis literature indicates a gap specifically around decision-making. There is adequate advice at a management level and a technical operational instructions level. But no adequate advice was found for decisions at a middle-level of granularity; specifically, gaps of note are:

- strategic selection of tactics, that is, which analysis heuristic or technical tool to employ in a particular situation and why
- when the investigator is justified in generalizing; that is, making a stronger, broader claim from singular pieces of evidence
- what information to report and how to communicate it to convince someone that the investigator should be believed

These problems are similar to those that scientists face in conducting their research. Incident responders have quite different operating environments than most scientists. However, as argued by Spring et al. [165], there is no obvious barrier to considering security research as a type of science. Scientific methods and norms need to be adapted, as security poses certain novel challenges. However, as argued by Spring and Illari [163], security analysts already take an approach to generalizing knowledge that is both similar to and can benefit from the wider literature on scientific explanation in philosophy of science. Therefore, it is plausible that CSIR would benefit from answering these gaps with methodology adapted from philosophy of science, though it is unlikely to be a panacea.

Various other disciplines might contribute to filling these decision-making gaps. Reviews of such other fields are future work. Note that none such arose through our review of CSIR standards. This gap indicates a reticence to take on scientific tools. The exception is the intelligence community, which genuinely integrates psychology and behavioral economics [79] and to some extent Bayesian belief propagation [27]. Other fields that might naturally be interrogated for links to incident analysis in future work include game theory, decision theory, information theory, systems engineering, internet measurement, and risk assessment. As one example, game theory and network security already have a developed overlap, for example see Alpcan and Başar [4], which might be adaptable to CSIR.

However, we are not aware of any work that has attempted to formalize decision-making in CSIR specifically. We leave this as an area of future work. The work by Horneman [82], adapting Heuer [79] to computer network analysis, is the closest approach so far. However, when our review calls for more structure to decision-making, it would mean further structure and specification of what she identifies as “analytical acumen” and its use.

Finally, we would like to note that this lack of published standards for this granularity of decision-making does not mean that incident responders are doing a bad job. Incident analysis is trade-craft, essentially handed down by apprenticeship, on-the-job training, and mentorship. FIRST has enabled such skill dissemination for decades. Likely, various norms of reasoning through the three gaps we note have been developed and disseminated amongst small groups of analysts. Getting access to and surveying these analysts and their reasoning methods would also be a rich area for future work, though it is fraught with difficulties of gaining trust, access, time, and representative samples. The abundance of CSIR teams, and the general social importance of responding to computer security incidents, has perhaps made this trade-craft approach unsustainable. The unsustainable nature of such trade-craft approaches heavily influences our belief that it is time to consolidate decision-making in incident analysis and begin publicly filling in these gaps.

List of Acronyms

ACM	Association for Computing Machinery
ACoD	Art into Science: A Conference for Defense
AES	Advanced Encryption Standard
AirCERT	Automated Incident Reporting
APWG	Anti-Phishing Working Group
ARMOR	Assistant for Randomized Monitoring Over Routes
ARPA	Advanced Research Projects Agency, from 1972–1993 and since 1996 called DARPA
ATT&CK	Adversarial Tactics, Techniques, and Common Knowledge (by MITRE)
BCP	Best Current Practice, a series of documents published by IETF
BGP	Border Gateway Protocol
BI	logic of bunched implications
BIS	Department for Business, Innovation, and Skills (United Kingdom)
BLP	Bell-Lapadula, a model of access control

CAE	Center of Academic Excellence
CAIDA	Center for Applied Internet Data Analysis, based at University of California San Diego
CAPEC	Common Attack Pattern Enumeration and Classification (by MITRE)
CCIPS	Computer Crime and Intellectual Property Section of the DoJ
CCE	Common Configuration Enumeration (by NIST)
CCSS	Common Configuration Scoring System (by NIST)
CEE	Common Event Expression (by MITRE)
CERT/CC	CERT® Coordination Center operated by Carnegie Mellon University
CIA	Central Intelligence Agency (US)
CIS	Center for Internet Security
CNA	Computer Network Attack
CND	Computer Network Defense
CNO	Computer Network Operations
CPE	Common Platform Enumeration (by NIST)
CSIR	Computer Security Incident Response
CSIRT	Computer Security Incident Response Team
CTL	Concurrent Time Logic
CVE	Common Vulnerabilities and Exposures (by MITRE)
CVRF	Common Vulnerability Reporting Framework
CVSS	Common Vulnerability Scoring System, maintained by FIRST
CWE	Common Weakness Enumeration (by MITRE)
CWSS	Common Weakness Scoring System, maintained by MITRE
CybOX	Cyber Observable Expression, maintained by MITRE
DARPA	Defense Advanced Research Projects Agency
DHS	US Department of Homeland Security
DNS	Domain Name System
DoD	US Department of Defense
DoJ	US Department of Justice
ENISA	EU Agency for Network and Information Security
EPSRC	Engineering and Physical Sciences Research Council (United Kingdom)
EU	European Union
FAA	Federal Aviation Administration (US)
FBI	US Federal Bureau of Investigation
FDA	US Food and Drug Administration
FIRST	Forum of Incident Response and Security Teams
FISMA	Federal Information Security Management Act (US)
FS-ISAC	Financial Services ISAC
GCHQ	Government Communications Headquarters (United Kingdom)
GFIRST	Government FIRST
HotSoS	Symposium on the Science of Security
HTTP	Hypertext Transfer Protocol, a standard by W3C
HTCIA	High Technology Crime Investigation Association
IC	intelligence community
ICT	information and communications technology
IEEE	Institute of Electrical and Electronic Engineers

IEP	Information Exchange Policy
IETF	Internet Engineering Task Force
IDS	intrusion detection system
IODEF	Incident Object Description Exchange Format
IODEF+	Incident Object Description Exchange Format Extensions (RFC 5901)
IDMEF	Intrusion Detection Message Exchange Format (RFC 4765)
ISAC	Information Sharing and Analysis Center
ISC	Internet Storm Centerpart of the privately run SANS
ISO	International Organization for Standardization
ISP	Internet Service Provider
ITU	International Telecommunications Union, an agency of the UN
LAX	Los Angeles International Airport
LBNL	Lawrence Berkeley National Laboratory
MAEC	Malware Attribute Enumeration and Characterization (by MITRE)
MITRE	the Mitre Corporation
MMDEF	Malware Metadata Exchange Format
MoD	Ministry of Defence (United Kingdom)
NATO	North Atlantic Treaty Organization
NCA	National Crime Agency (UK)
NCCIC	US National Cybersecurity and Communications Integration Center
NDA	non-disclosure agreement
NIDPS	Network Intrusion Detection and Prevention System
NIST	National Institute of Standards and Technology, part of the US Department of Commerce
NSA	National Security Agency (US)
NSF	National Science Foundation (US)
OCIL	Open Checklist Interactive Language (by NIST)
OVAL	Open Vulnerability and Assessment Language (by MITRE)
OWASP	Open Web Application Security Project
OWL	Ontology Web Language
pDNS	passive DNS traffic analysis
PSIRT	Product Security Incident Response Team
RAM	Random Access Memory
RCT	Randomized Controlled Trial
REN-ISAC	Research and Education Networking ISAC
RID	Real-time Inter-network Defense
RISCS	Research Institute in Science of Cyber Security (United Kingdom)
RFC	Request for Comments, standardization and informational documents published by the IETF
SANS Institute	Sysadmin, Audit, Network, and Security Institute
SCAP	Security Content Automation Protocol (by NIST)
SiLK	System for Internet-level Knowledge, an open-source analysis tool set published by CERT/CC
SoK	Systematization of Knowledge paper in IEEE Oakland conference
STIX	Structured Threat Information Expression (by MITRE)

STS	Science and Technology Studies (a field synthesizing philosophy of science, history of science, sociology of science, and philosohpy of technology)
TAXII	Trusted Automated eXchange of Indicator Information (by MITRE)
TCP/IP	Transmission Control Protocol / Internet Protocol
TLA	Temporal Logic of Actions
TLP	Traffic Light Protocol
TLD	Top-Level Domain (in DNS)
TSA	Transport Security Administration (US)
TTPs	Tools, tactics, and procedures
UN	United Nations
US	United States of America
US-CERT	US Computer Emergency Readiness Team, a branch of NCCIC within DHS
URL	Uniform Resource Locator
VERIS	Vocabulary for Event Recording and Incident Sharing
W3C	World Wide Web Consortium
XCCDF	Extensible Configuration Checklist Description Format (by NIST)
XML	Extensible Markup Language, a standard by W3C

APPENDICES

A CLARIFICATIONS ON CHOICES MADE FOR TABLE 11

In this Appendix, we comment further on papers in Table 11, explaining how we categorize them, and commenting on some not further discussed in Section 5.2 above. We include comment on the ISO 27000 standards generally, as well.

NIST SP 800-83 [160] Section 4.2 is titled "detection and analysis," yet we have provocatively labeled the document as having no bearing on the analysis phase. The section's advice on analysis is entirely tool-focused pragmatics. Analysis should take place on an isolated or virtualized operating system to prevent spread of infection, and so on. The document mentions some fields that the investigator may want to collect, such as file names, service ports, and "how to remove the malware." There is no advice on how to obtain this information, why, or what it might be useful for. Therefore, these are best understood as reporting constraints, not analysis advice. This result is disappointing, considering Section 4 gives its opening motivation as "this section of the guide builds on the concepts of SP 800-61 by providing additional details about responding to malware incidents."

NIST SP 800-86 Kent et al. [102] suffers similarly to SP 800-83; it consists of a stream of data formats and types and assumes that the investigator will know what to do now that the possible data types have been listed. These make up underlying technical skills necessary for an investigation, and so are not completely irrelevant to incident response. However, they do not help us understand how investigators make decisions. The extent of the advice on analysis again amounts to essentially reporting and collection constraints, respectively: "the analysis should include identifying people, places, items, and events, and determining how these elements are related... often, this effort will include correlating data among multiple sources" [102, pp. 3–6].

The ISO 27000-series is dedicated to information security. We have identified five standards that are within our scope of the particular parts of incident response. The relationship between these standards is documented in each of the standards by what is consistently labeled Figure 1 there. That figure states clearly that all the listed standards are applicable to "investigation process classes and activities" [90, p. ix]. Process classes are readiness, initialization, acquisitive, and investigative; activities overlap these classes, and are plan, prepare, respond, identify-collect-acquire-preserve, understand, report, and close. This taxonomy is essentially consistent with the taxonomies used by the IETF, NIST, and FIRST [132].

However, where a NIST standard such as SP 800-61 is a single 70-page document, the ISO incident response standards are each 10–15 pages of unique content with 10–15 pages that are repeated in each document. Thus, the five ISO documents combined are comparable in scope and detail to SP 800-61. However, unlike a NIST publication, the ISO documents do not present a clear investigative goal among the options we have distinguished, even within documents, let alone among them.

27043, for example, seems to unknowingly alternate between incident response for fixing systems and analysis for providing evidence to a legal proceeding. Within 27043 ISO/IEC [90], Section 8 reads like advice from CERT/CC [2], and Section 9 reads like advice from Casey [33, ch. 2]. The shift is abrupt and without explanation. The shift includes a shift in terminology and jargon for referring to essentially the same mental process by the investigator. This oddity does not build confidence that the ISO standards actually present a unified methodology for incident response as a series of disconnected vignettes.

27041 does little to dispel this sense of disconnectedness. This ISO document is disconnected from the other incident management documents in that it focuses on the client-contractor relationship. The sense in which is a process is validated is that "the work instruction britishfulfils the requirements agreed with the client" [88, p. 9]. 27041 ISO/IEC [88, pp. 12–13] states an investigation composed of validated examinations "can be considered to be validated" while defining a validated examination as one mode up of validated processes. Assuming composability of valid processes is a dangerous claim. Concurrent program verification has shown such claims cannot be assumed and are challenging to prove [125, 134]; albeit the technical sense of "valid" is slightly different, doubt in the ISO assumption seems warranted.

For these reasons, the ISO standards would struggle to function well as a unified whole. There does not seem to be an overarching editorial guidance to assure consistency or navigate conflicts. At best, if the reader already knows how to navigate the different, conflicting contexts, the ISO documents are useful expressions of each area of concern. The level of detail is appropriate for ensuring management ability to oversee a process, rather than to do the process itself. Even the most specific documents (27037 and 27042), to which the other, more general documents refer for details, are thin on anything that might help with actual decision-making.

RFC 2196 is quite old, and its advice shows its age. The steps are in general sound; however, they are from a time when it was reasonable to ask for "all system events (audit records)" to be recorded and evaluated by the investigator [61, p. 54]. The text assumes that incident investigators will know what to do with these events once logged. This advice is not bad, such as it is; however, it is best understood as historical rather than actionable advice.

RFC 6545 [129] and RFC 6546 jointly detail Real-time Inter-network Defense. RFC 6545 describes conceptual and formal details, whereas RFC 6546 provides technical communication and encryption details. RFC 6545 is an extension of IODEF [46], specifying methods and norms of communication using IODEF between organizations. As such, the document focuses on what to report, and how to use reports for mitigation. Policy of use and sensitivity of information is explicitly integrated into the format. How analysis produces adequate data is out of scope. However, by providing such explicit standards on what should be reported and how those reports can expect to be used, RFC 6545 does put constraints on analysis and evidence collection—those phases need to produce reporting with the specified types of fields.

RFC 7203 extends IODEF [46] "to embed and convey various types of structured information" Takahashi et al. [173, p. 2]. Specifically, the various metrics and formats such as CVSS and CVE captured in Table 6. This extension serves to integrate two types of reporting format and constraint. This is useful, but is mostly programmatic. Therefore, it is not directly about reporting in the same way RFC 7970 is. Although technically detailed, from a decision-making point of view RFC 7203 just suggests that these metrics are useful ways to describe an incident and report on it, and that investigators should do so. RFC 5901 makes similar suggestions specifically for reporting phishing [24].

RFC 8134 is informational, and not a standard. It provides a list of information exchanges, collaborations, and implementations that make use of IODEF. Because information exchanges are a source of evidence collection, the details about what information is available from what groups provides evidence collection suggestions and introductions. Although this advice at a rather abstract level, it is useful, because it provides a discussion of network defense information sharing arrangements that is not commonly quite so public.

Gorzelak et al. [66] is a study commissioned by ENISA and executed by the Polish CERT. The focus is on data sources—how do CSIRTs monitor their constituents. The method employed is a survey of over 100 CSIRTs. While this data is at best instructive of where to get data, it is an important resource for how respondents evaluate the quality of data sources. Such evaluation is directly relevant to evidence collection decisions. It is unlikely this study is instructive outside this relatively narrow context. However, it is directly relevant context for this work.

Alberts et al. [2] is primarily about contextualizing incident management within a wider organizational context. In fact, Alberts et al. [2, pp. 24–26] is one of the best assessments of the relationship of investigation to preparation and protection we have found in this review. However, our focus is not on situation of the investigative process within an organization. Alberts et al. [2, p. 128ff.] is an ambitious effort to organize a flow chart for incident response. Because their scope includes technical, management, and legal responses, the level of detail devoted to analysis amounts to "Designated personnel analyze each event and plan, coordinate, and execute the appropriate technical response" [2, p. 136].

B LIMITATIONS

While our methods have much to recommend them, there are of course limitations. Some of these are practical, such as the restriction to publications available in English. Some limitations are a function of restricting the scope to standards. Another concern is the representativeness or generalizability from rather old case studies of incident analysis. Perhaps the most dangerous limitation is a result of the subject matter—security practitioners tend to be secretive about their methodologies. We discuss each of these limitations in turn.

The restriction to English will naturally limit the results. For example, an internet search for "信息安全事件应对" (information security incident response) returns a couple dozen results on Google as well as Google Scholar. This seems to be the preferred term in mainland China. A search for "电脑安全事件应对" (computer security incident response) returns only a couple of Taiwanese sites.

The importance of this language choice on actually limiting documents available to us is less clear. EU and UN documents would be available in English as a matter of policy. The US government, which publishes in English, dominates in this space, as do US companies. Countries that are not allied to the US and have developed computer security capabilities are relatively few; basically just the Russian Federation and the **People's Republic of China (PRC)**. While it is possible these countries have published comprehensive decision-making procedures for incident response, it seems unlikely. It also seems likely that, given how much attention the US security establishment pays to Russia and the PRC, if such a thing were published it would be found and reported on, if not translated. For these reasons, we judge the impact of limiting our search to English documents is a low risk.

Focusing on standards, and what they cite, keeps the scope manageable but also creates other limitations. The type of information published in standards is different than that in academic journals and conferences, and this imposes some limits on the work. Specifically, standards are on a slower publication cycle than academic work. This delay would be a problem if our topic were extensively covered in the academic literature. However, as indicated by Section 3, academic publications do not appear to cover decision-making during computer security incident response.

Creation of technology standards is itself a complex process, and as Section 2.2 touches on, the process has a complex history in its own right [166]. The way standards are made imposes its own limitations on our findings. Standards are rarely made purely for the dissemination of information; rather, they usually solidify a dominant business position. Security standards appear to be an outlier from this norm, as they have unique concerns about correctness and non-subversion [106]. Incident response standards are essentially unstudied within the academic standards literature; Kuhlmann et al. [106] mostly focus on cryptographic standards. This situation means we accept a risk in that we do not know what biases may be embedded in the creation of incident response standards. The standards literature provides evidence there will be a bias, but has not studied incident response standards to provide evidence for what that bias might be. We would like to know whose interests are best served by the creation of incident response standards, for example.

Case studies or collections of cases that have been analyzed by others provide demonstrations of what sort of attacks are possible. We have two of the earliest examples of this style of reporting with Stoll [168] and Cheswick [35]. It is important to note these are examples. A survey of incident case studies may be a useful additional project, though it is out of our scope. The practitioners who wrote these standards documents would be aware that many security vendors publish accounts of adversaries they have tracked in the course of their work. These are of varying quality, scope, and importance. More recent impactful studies include, for example, Mandiant tracking an alleged unit of the Chinese military [119] and Google's self-report of compromise attributed to China [48]. Some case reports are official government commissioned, such as the Black Tulip report analyzing the compromise of a TLS certificate authority critical to the operation of many Dutch government websites [81].

The scope need not be an individual case. Some studies focus on trends rather than individual cases. Verizon's Data Breach Investigation Report is probably the best-known example (see, e.g., Verizon [183], ,184]). United States federal civilian agencies must make annual breach reports to Congress per FISMA requirements; such detailed reports have been examined for trends [140].

A closely related limitation of concern involves secrecy. Many incident response organizations may not wish to disclose their processes and procedures in detail, lest the adversary learn how to subvert or avoid them. Other areas of information security experience similar publication restrictions [163]. Incident responders likely have a legitimate concern in this regard, and may also have a legal or regulatory requirement to keep certain information or processes private. Therefore, this limitation imposes a significant risk that relevant information is not public. Lack of access obviously limits the review. Our approach to reduce the impact of this limitation is to understand that we ought to read between the lines of available documents when we can justify expanding our interpretation of a document's contents with circumstantial evidence from the context surrounding its publication. However, we must accept that there is an amount of information about our topic that simply is not public and we cannot hope to access for a public literature review. One could perhaps use news articles and audit reports to attempt to evidence the extent to which organizations in fact implement the available standards; we leave such investigations for future work.

C EXTRA REFERENCES PROVIDED BY REVIEWERS

In two related but distinct areas, our anonymous reviewers highlighted some important context for our review of CSIR standards. The first element is that norms or standards for CSIR are not separable from those for human decision-making in digital forensics. There is certainly some overlap; ISO in particular (Table 7) references more work from the digital forensics community than the other standards surveyed in Section 4. The unresolved question is why these documents are not cited by the CSIR standards if they are understood to be the expected reasoning methodology.

A related but distinct thread of critique was that there are some studies of CSIRTs in practice we did not discuss. We mentioned the anthropological work of Sundaramurthy et al. [171], but we have not offered a survey of studies of CSIR practitioners. Section 3 suggests that there is no ready survey of such work, either. Similarly, if these surveys of practitioner behavior have prompted reflection and influenced practitioners, then there is an open question as to why they are not referenced by the standards.

The reviewers' experience suggests there is some tacit knowledge among CSIR professionals that they would be well-served to read the digital forensics literature. However, one reviewer commented that "[w]hile intrusion

Table 12. Categorization of New Reviewer-supplied Documents

	Directness per Phase			Scope per Goal					
Document	Col	Anz	Rep	Fix	Int	LE	General	Type	Formal
Ahmad et al. [1]	╳	╳	╳	B	╳	╳	Un	Study	Qual
Böhme et al. [18]	C	╳	╳	╳	╳	M	Un	Ont	∅
Casey [32]	C	C	╳	N	╳	N	Un	Study	∅
Chow et al. [39]	D	╳	╳	N	N	N	Likely	Study	∅
Farmer and Venema [56]	D	╳	╳	N	N	N	High	Study	∅
Grispos et al. [68]	C	╳	╳	B	╳	B	Likely	Study	∅
Grispos et al. [69]	╳	╳	╳	B	╳	╳	Un	Study	∅
Grispos et al. [70]	╳	╳	C	M	╳	╳	Un	Study	∅
He et al. [77]	╳	╳	D	M	╳	╳	Likely	Study	Qual
Hove et al. [83]	╳	╳	╳	B	╳	╳	Un	Study	∅
Luttgens et al. [117]	D	C	D	B	╳	N	High	Adv	∅
Orderløkken [136]	╳	╳	╳	B	╳	╳	Un	Study	∅
Rollason-Reese [148]	C	C	C	B	╳	╳	Un	Study	Qual
Shedden et al. [155]	╳	╳	╳	B	╳	╳	Un	Study	∅
Vangelos [182]	╳	╳	C	B	╳	╳	Un	Ont	Perf
Werlinger et al. [186]	╳	╳	╳	B	╳	╳	Un	Study	Perf

The phases are collection, analysis, and reporting. Cells have a cross if a phase is not addressed. Advice directness values are direct (D) or constraints-based (C). Values for goals are to fix an infected system (fix), gathering intelligence (int), and law enforcement (LE) action. Values for a document's intended scope are narrow (N), medium (M), or broad (B). Values for generalizability of an approach are unlikely (Un), likely (Likely), highly likely (High), or already widely generalizable (Wide). Document types are case studies (Study), ontologies (Ont), advice on actions (Adv), and explicit instructions (Instr). Values for formalization are not present (∅), qualitative (Qual), formal, or perfunctory (Perf).

detection and incident response work together rather often, real forensics [practitioners] are not [seen] that often at incident responder conference[s] (maybe it is too academic?)." This situation indicates a tension between what experienced CSIR professionals expect and the information to which a junior CSIR professional has public access.

It is not clear whether the problem lies in the CSIR standards themselves or our search methodology choice. As far as documenting what the standards contain, our search and analysis methodology is sound; or at least, no reviewer raised material doubts that it accurately captures the state of CSIR standards. Section 2 documented various ways in which the standards literature is only an approximation of what CSIR practitioners actually do. We believe the reviewers' additional suggested documents are a consequence of a mismatch between standards and practice. We do not provide a detailed analysis of each of these 17 papers, but we provide a summary analysis using the same categorization criteria as used in Table 11.

Table 12 captures our analysis of the documents suggested by the reviewers. Five documents (Böhme et al. [18], Casey [33], Chow et al. [39], Farmer and Venema [56], and Luttgens et al. [117]) were provided as examples of documents from digital forensics used in educating CSIR practitioners. As Casey [33] was already captured by our review process and is summarized in Section 5, it is not included in Table 12. The other 12 documents in Table 12 are studies of CSIRTs *in situ*. Conducting security-sensitive human-centric research has various challenges [105] that these *in situ* studies do not consistently address. In general, these works focus on how to enable CSIR practitioners to learn on the job, how a CSIRT can be managed, and technical details of how computers work. Our review presented in this article asks about what practitioners should learn to connect such technical knowledge. These threads of work are related, but none of these documents fill in the gaps in publicly available advice to CSIR practitioners identified by our review. Anything with three crosses in the phases columns of Table 12 would have failed our initial filtering methodology described in Section 4.

Fifteen of the 16 suggested documents are not directly related to our research question. This provides some evidence that our search methodology provided as good a view of recommended CSIR practice as is publicly available. Some of the documents cite and build on standards we have surveyed above, but without the benefit of the broader context of the landscape of competing standards in CSIR that we have provided. Some of the documents conduct a limited case study and propose that other CSIR professionals should learn from it; however, such studies represented in Table 12 are uncontextualized and have little detail about analyst thought processes. More often than not, they are about management's relation to the CSIRT than CSIR practice. Our hypothesis for why our review finds they are not cited by existing standards is that they have not, in fact, influenced CSIR practice.

Luttgens et al. [117] is the exception. The book comes close to meeting our goals and filling in many gaps in advice to CSIR professionals. This is why we now discuss the book in Section 5.

REFERENCES

[1] Atif Ahmad, Justin Hadgkiss, and Anthonie B. Ruighaver. 2012. Incident response teams–challenges in supporting the organisational security function. *Comput. Sec.* 31, 5 (2012), 643–652.

[2] Chris Alberts, Audrey Dorofee, Georgia Killcrece, Robin Ruefle, and Mark Zajicek. 2004. Defining incident management processes for CSIRTs: A work in progress. Technical Report CMU/SEI-2004-TR-015, Software Engineering Institute, Carnegie Mellon University.

[3] Christopher Alberts, Audrey Dorofee, Robin Ruefle, and Mark Zajicek. 2014. An introduction to the mission risk diagnostic for incident management capabilities (MRD-IMC). Technical Report CMU/SEI-2014-TN-005, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA.

[4] Tansu Alpcan and Tamer Başar. 2011. *Network Security: A Decision and Game-theoretic Approach.* Cambridge University Press, Cambridge, U.K.

[5] Bernhard Amann, Robin Sommer, Aashish Sharma, and Seth Hall. 2012. A lone wolf no more: Supporting network intrusion detection with real-time intelligence. In *Research in Attacks, Intrusions, and Defenses.* Springer, 314–333.

[6] Ross Anderson, Chris Barton, Rainer Böhme, Richard Clayton, Michel J. G. Van Eeten, Michael Levi, Tyler Moore, and Stefan Savage. Measuring the cost of cybercrime. In *Workshop on the Economics of Information Security.*

[7] M. Antonakakis, R. Perdisci, W. Lee, N. Vasiloglou II, and D. Dagon. 2011. Detecting malware domains at the upper DNS hierarchy. In *20th Usenix Security Symposium.*

[8] Sasikanth Avancha, Amit Baxi, and David Kotz. 2012. Privacy in mobile technology for personal healthcare. *ACM Comput. Surv.* 45, 1 (2012), 3:1–3:54.

[9] Franz Baader. 2003. *The Description Logic Handbook: Theory, Implementation and Applications.* Cambridge University Press, Cambridge, U.K.

[10] Rebecca Bace and Peter Mell. 2001. Intrusion detection systems. Technical Report SP 800-31, U.S. National Institute of Standards and Technology, Gaithersburg, MD.

[11] Susan Ballou, Jaime Carazo, Bill Crane, Fred Demma, Grant Gottfried, Sam Guttman, Jeffry Herig, Tim Hutchison, David Icove, Bob Jarzen, Tom Johnson, Karen Matthews, Mark Pollitt, David Poole, Mary Riley, Kurt Schmid, Howard A. Schmidt, Raemarie Schmidt, Carl Selavka, Steve Sepulveda, Todd Shipley, Chris Stippich, Carrie Morgan Whitcomb, and Wayne Williams. 2001. *Electronic Crime Scene Investigation: A Guide for First Responders.* U.S. Department of Justice, Office of Justice Programs, National Institute of Justice, Washington, DC.

[12] Konstantia Barmpatsalou, Tiago Cruz, Edmundo Monteiro, and Paulo Simoes. 2018. Current and future trends in mobile device forensics: A survey. *ACM Comput. Surv.* 51, 3 (2018), 46:1–46:31. DOI: http://doi.acm.org/10.1145/3177847

[13] Jon Barwise and Jerry Seligman. 1997. *Information Flow: The Logic of Distributed Systems.* Cambridge University Press.

[14] Nicole Lang Beebe and Jan Guynes Clark. 2005. A hierarchical, objectives-based framework for the digital investigations process. *Dig. Invest.* 2, 2 (2005), 147–167.

[15] Steve Bellovin. 1992. There be dragons. In *USENIX Security Symposium.*

[16] Vilius Benetis, Olivier Caleff, Cristine Hoepers, Angela Horneman, Allen Householder, Klaus-Peter Kossakowski, Art Manion, Amanda Mullens, Samuel Perl, Daniel Roethlisberger, Sigitas Rokas, Mary Rossell, Robin M. Ruefle, Desiree Sacher, Krassimir T. Tzvetanov, and Mark Zajicek. 2019. Computer security incident response team (CSIRT) services framework. Technical Report ver. 2.1, FIRST, Cary, NC.

[17] Robert Biddle, Sonia Chiasson, and P. C. Van Oorschot. 2012. Graphical passwords: Learning from the first twelve years. *ACM Comput. Surv.* 44, 4 (2012), 19:1–19:41.

[18] Rainer Böhme, Felix C. Freiling, Thomas Gloe, and Matthias Kirchner. 2009. Multimedia forensics is not computer forensics. In *International Workshop on Computational Forensics.* Springer, LNCS 5718, 90–103.

[19] Marcus Botacin, Paulo Lício De Geus, and André Grégio. 2018. Who watches the watchmen: A security-focused review on current state-of-the-art techniques, tools, and methods for systems and binary analysis on modern platforms. *ACM Comput. Surv.* 51, 4 (2018), 69:1–69:34. DOI: http://doi.acm.org/10.1145/3199673

[20] S. L. Brand. 1985. Dod 5200.28-std Department of Defense trusted computer system evaluation criteria (orange book). Technical Report. US Department of Defense.

[21] Susan W. Brenner. 2002. Organized cybercrime: How cyberspace may affect the structure of criminal relationships. *North Carol. J. Law Technol.* 4 (2002), 1.

[22] Charles Brookson, Scott Cadzow, Ralph Eckmaier, Jörg Eschweiler, Berthold Gerber, Alessandro Guarino, Kai Rannenberg, Jon Shamah, and Sławomir Górniak. 2015. Definition of cybersecurity: Gaps and overlaps in standardisation. Technical Report v1.0, ENISA, Heraklion, Greece.

[23] N. Brownlee and E. Guttman. 1998. Expectations for Computer Security Incident Response. RFC 2350 (Best Current Practice). https://tools.ietf.org/html/rfc2350.

[24] P. Cain and D. Jevans. 2010. Extensions to the IODEF-Document Class for Reporting Phishing. RFC 5901 (Proposed Standard). https://tools.ietf.org/html/rfc5901.

[25] Sergio Caltagirone. 2005. Evolving active defense strategies. Technical Report CSDS-DF-TR-05-27, University of Idaho, Moscow, ID.

[26] Sergio Caltagirone and Deborah Frincke. 2005. Adam: Active defense algorithm and model. *Aggressive Network Self-Defense*. O'Reilly Media, Inc. 287–311.

[27] Sergio Caltagirone, Andrew Pendergast, and Christopher Betz. 2013. The diamond model of intrusion analysis. Technical Report, Center for Cyber Intelligence Analysis and Threat Research. Retrieved from http://www.threatconnect.com/methodology/diamond_model_of_intrusion_analysis.

[28] Stefano Calzavara, Riccardo Focardi, Marco Squarcina, and Mauro Tempesta. 2017. Surviving the web: A journey into web session security. *ACM Comput. Surv.* 50, 1 (2017), 13:1–13:34.

[29] Richard Caralli, James Stevens, Lisa Young, and William Wilson. 2007. Introducing OCTAVE Allegro: Improving the information security risk assessment process. Technical Report CMU/SEI-2007-TR-012, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA.

[30] Brian Carrier and Eugene H. Spafford. 2003. Getting physical with the digital investigation process. *Int. J. Dig. Evid.* 2, 2 (2003), 1–20.

[31] Brian Carrier and Eugene H. Spafford. 2004. An event-based digital forensic investigation framework. In *the Digital Forensic Research Workshop*.

[32] Eoghan Casey. 2005. Case study: Network intrusion investigation–lessons in forensic preparation. *Dig. Invest.* 2, 4 (2005), 254–260.

[33] Eoghan Casey. 2010. *Handbook of Digital Forensics and Investigation*. Elsevier.

[34] Varun Chandola, Arindam Banerjee, and Vipin Kumar. 2009. Anomaly detection: A survey. *ACM Comput. Surv.* 41, 3 (2009), 15:1–15:58.

[35] Bill Cheswick. 1992. An evening with Berferd: In which a cracker is lured, endured, and studied. In *USENIX Winter Technical Conference*.

[36] W. R. Cheswick, S. M. Bellovin, and A. D. Rubin. 2003. *Firewalls and Internet Security: Repelling the Wily Hacker* (2nd ed.). Addison-Wesley Professional.

[37] Tung Ju Chiang, Jen Shiang Kouh, and Ray-I. Chang. 2009. Ontology-based risk control for the incident management. *Int. J. Comput. Sci. Netw. Secur.* 9, 11 (2009), 181.

[38] Jin-Hee Cho, Shouhuai Xu, Patrick M. Hurley, Matthew Mackay, Trevor Benjamin, and Mark Beaumont. 2019. STRAM: Measuring the trustworthiness of computer-based systems. *ACM Comput. Surv.* 51, 6 (2019), 128:1–128:47. DOI: http://doi.acm.org/10.1145/3277666

[39] Jim Chow, Ben Pfaff, Tal Garfinkel, and Mendel Rosenblum. 2005. Shredding your garbage: Reducing data lifetime through secure deallocation. In *USENIX Security Symposium*. 22–22.

[40] Séamus Ó Ciardhuáin. 2004. An extended model of cybercrime investigations. *Int. J. Dig. Evid.* 3, 1 (2004), 1–22.

[41] Paul Cichonski, Tom Millar, Tim Grance, and Karen Scarfone. 2012. Computer security incident handling guide. Technical Report SP 800-61r2, U.S. Dept. of Commerce, National Institute of Standards and Technology, Gaithersburg, MD.

[42] Fred Cohen, Julie Lowrie, and Charles Preston. 2011. The state of the science of digital evidence examination. In *theInternational Federation for Information Processing Conference (IFIP'11).*

[43] Frederick B. Cohen. 1995. *Protection and Security on the Information Superhighway*. John Wiley & Sons, Inc.

[44] Frederick B. Cohen. 2010. *Fundamentals of Digital Forensic Evidence*. Springer, New York, 790–808.

[45] Andrew Cormack. 2015. Janet suggested charter for system administrators. Technical Report, Jisc, Bristol, U.K.

[46] R. Danyliw. 2016. The Incident Object Description Exchange Format Version 2. RFC 7970 (Proposed Standard). https://tools.ietf.org/html/rfc7970.

[47] R. Danyliw, J. Meijer, and Y. Demchenko. 2007. The Incident Object Description Exchange Format. RFC 5070 (Proposed Standard). Retrieved from https://www.rfc-editor.org/rfc/rfc5070.txt.

[48] David Drummond. 2010. A new approach to China. Retrieved from http://googleblog.blogspot.com/2010/01/new-approach-to-china.html.

[49] Felicia Duran, Stephen H. Conrad, Gregory N. Conrad, David P. Duggan, and Edward Bruce Held. 2009. Building a system for insider security. *IEEE Secur. Priv.* 7, 6 (2009), 30–38.

[50] Harry T. Edwards, Constantine Gatsonis, Margaret A. Berger, Joe S. Cecil, M. Bonner Denton, Marcella F. Fierro, Karen Kafadar, Pete M. Marone, Geoffrey S. Mearns, Randall S. Murch, Channing Robertson, Marvin E. Schechter, Robert Shaler, Jay A. Siegel, Sargur N. Srihari, Sheldon M. Wiederhorn, and Ross E. Zumwalt. 2009. *Strengthening Forensic Science in the United States: A Path Forward.* National Academies Press, Washington, DC.

[51] Matthew Edwards, Awais Rashid, and Paul Rayson. 2015. A systematic survey of online data mining technology intended for law enforcement. *ACM Comput. Surv.* 48, 1 (2015), 15:1–15:54.

[52] Manuel Egele, Theodoor Scholte, Engin Kirda, and Christopher Kruegel. 2008. A survey on automated dynamic malware-analysis techniques and tools. *ACM Comput. Surv.* 44, 2 (2008), 6:1–6:42.

[53] Andreas Ekelhart, Stefan Fenz, Markus Klemen, and Edgar Weippl. 2007. Security ontologies: Improving quantitative risk analysis. In *Hawaii International Conference on System Sciences.*

[54] ENISA. 2006. A step-by-step approach on how to set up a CSIRT. Technical Report WP2006/5.1, Heraklion, Greece. European Union Agency for Cybersecurity.

[55] ETSI. 2014. Key performance security indicators (KPSI) to evaluate the maturity of security event detection. Technical Report GS ISI 003 V1.1.2, ETSI Information Security Indicators (ISI), Cedex, France.

[56] Dan Farmer and Wietse Venema. 2009. *Forensic Discovery.* Addison-Wesley Professional.

[57] Stefan Fenz and Andreas Ekelhart. 2009. Formalizing information security knowledge. In *Symposium on Information, Computer, and Communications Security.* ACM, 183–194.

[58] P. Ferguson and D. Senie. 2000. Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing. RFC 2827 (Best Current Practice) Retrieved from https://www.rfc-editor.org/rfc/rfc2827.txt.

[59] FIRST. FIRST vision and mission statement. Retrieved from https://first.org/about/mission.

[60] FIRST. 2017. Security Reference Index. Retrieved from https://first.org/resources/guides/reference.

[61] B. Fraser. 1997. Site Security Handbook. RFC 2196 (Informational). https://tools.ietf.org/html/rfc2196.

[62] Peter Galison. 2012. Augustinian and Manichaean science. In *Symposium on the Science of Security.*

[63] Simson Garfinkel, Paul Farrell, Vassil Roussev, and George Dinolt. 2009. Bringing science to digital forensics with standardized forensic corpora. *Dig Invest.* 6 (2009), S2–S11.

[64] GÉANT. 2020. About GÉANT. Retrieved from www.geant.org/About.

[65] GÉANT. 2020. Transits-i course materials. Retrieved from www.geant.org/Services/Trust_identity_and_security/Pages/TRANSITS-course-materials.aspx.

[66] Katarzyna Gorzelak, Tomasz Grudziecki, Paweł Jacewicz, Przemysław Jaroszewski, Łukasz Juszczyk, and Piotr Kijewski. 2011. Proactive detection of network security incidents. Technical Report 2011-12-07, CERT Polska/NASK, Heraklion, Greece.

[67] Tim Grance, Tamara Nolan, Kristin Burke, Rich Dudley, Gregory White, and Travis Good. 2006. Guide to test, training, and exercise programs for it plans and capabilities. Technical Report SP 800-84, U.S. Dept. of Commerce, National Institute of Standards and Technology, Gaithersburg, MD.

[68] George Grispos, William Bradley Glisson, and Tim Storer. 2015. Security incident response criteria: A practitioner's perspective. *arXiv preprint arXiv:1508.02526*, 2015.

[69] George Grispos, William Bradley Glisson, David Bourrie, Tim Storer, and Stacy Miller. 2017. Security incident recognition and reporting (SIRR): An industrial perspective. In *Americas Conference on Information Systems.*

[70] George Grispos, William Glisson, and Tim Storer. 2019. How good is your data? Investigating the quality of data generated during security incident response investigations. In *Hawaii International Conference on System Sciences.*

[71] E. Guttman, L. Leong, and G. Malkin. 1999. Users' Security Handbook. RFC 2504 (Informational). https://tools.ietf.org/html/rfc2504.

[72] Katie Hafner and Matthew Lyon. 1998. *Where Wizards Stay Up Late: The Origins of the Internet.* Simon and Schuster.

[73] Ryan Hankins, Tetsutaroh Uehara, and Jigang Liu. 2009. A comparative study of forensic science and computer forensics. In *Secure Software Integration and Reliability Improvement.* IEEE, 230–239.

[74] Eric Hatleback and Jonathan M. Spring. 2014. Exploring a mechanistic approach to experimentation in computing. *Philos. Technol.* 27 (3), 441–459.

[75] Douglas M. Hawkins. 2004. The problem of overfitting. *J. Chem. Inf. Comput. Sci.* 44, 1 (2004), 1–12.

[76] L. T. C. Ashton Hayes. 2008. Defending against the unknown: Antiterrorism and the terrorist planning cycle. *The Guardian* 10, 1 (2008), 32–36.

[77] Ying He, Chris Johnson, Karen Renaud, Yu Lu, and Salem Jebriel. 2014. An empirical study on the use of the generic security template for structuring the lessons from information security incidents. In *the 6th International Conference on Computer Science and Information Technology (CSIT'14).* IEEE, 178–188.

[78] Cormac Herley and P. C. van Oorschot. 2017. SoK: Science, security, and the elusive goal of security as a scientific pursuit. In *Symposium on Security and Privacy.* IEEE.

[79] Richards J. Heuer, Jr. 1999. *Psychology of Intelligence Analysis.* U.S. Central Intelligence Agency.

[80] Ivan Homoliak, Flavio Toffalini, Juan Guarnizo, Yuval Elovici, and Martín Ochoa. 2019. Insight into insiders and it: A survey of insider threat taxonomies, analysis, modeling, and countermeasures. *ACM Comput. Surv.* 52, 2 (2019), 30:1–30:40. DOI: http://doi.acm.org/10.1145/3303771

[81] Hans Hoogstraaten. 2012. Black Tulip: Report of the investigation into the DigiNotar Certificate Authority breach. Technical Report, Fox-IT.

[82] Angela Horneman. 2017. How to think like an analyst. Retrieved from https://insights.sei.cmu.edu/sei_blog/2017/07/how-to-think-like-an-analyst.html.

[83] Cathrine Hove, Marte Tårnes, Maria B. Line, and Karin Bernsmed. 2014. Information security incident management: Identified practice in large organizations. In *IT Security Incident Management & IT Forensics*. IEEE, 27–46.

[84] Eric M. Hutchins, Michael J. Cloppert, and Rohan M. Amin. 2011. Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains. *Lead. Issues Inf. Warf. Sec. Res.* 1 (2011), 80.

[85] C. Inacio and D. Miyamoto. 2017. Management Incident Lightweight Exchange (MILE) Implementation Report. RFC 8134 (Informational). https://tools.ietf.org/html/rfc8134.

[86] Chadni Islam, Muhammad Ali Babar, and Surya Nepal. 2019. A multi-vocal review of security orchestration. *ACM Comput. Surv.* 52, 2 (2019), 37:1–37:45. DOI: http://doi.acm.org/10.1145/3305268

[87] ISO/IEC. 2012. Information technology – security techniques – guidelines for identification, collection, acquisition and preservation of digital evidence. Technical Report 27037:2012, International Organization for Standardization and International Electrotechnical Commission, Geneva.

[88] ISO/IEC. 2015. Information technology – security techniques – guidance on assuring suitability and adequacy of incident investigative method. Technical Report 27041:2015, International Organization for Standardization and International Electrotechnical Commission, Geneva.

[89] ISO/IEC. 2015. Information technology – security techniques – guidelines for the analysis and interpretation of digital evidence. Technical Report 27042:2015, International Organization for Standardization and International Electrotechnical Commission, Geneva.

[90] ISO/IEC. 2015. Information technology – security techniques – incident investigation principles and processes. Technical Report 27043:2015, International Organization for Standardization and International Electrotechnical Commission, Geneva.

[91] ISO/IEC. 2016. Information technology – security techniques – information security incident management – part 1: Principles of incident management. Technical Report 27035-1:2016, International Organization for Standardization and International Electrotechnical Commission, Geneva.

[92] Mohammad Hanif Jhaveri, Orcun Cetin, Carlos Gañán, Tyler Moore, and Michel Van Eeten. 2017. Abuse reporting and the fight against cybercrime. *ACM Comput. Surv.* 49, 4 (2017), 68:1–68:27.

[93] Wenjun Jiang, Guojun Wang, Md Zakirul Alam Bhuiyan, and Jie Wu. 2016. Understanding graph-based trust evaluation in online social networks: Methodologies and challenges. *ACM Comput. Surv.* 49, 1 (2016), 10:1–10:35.

[94] Wolfgang John and Tomas Olovsson. 2008. Detection of malicious traffic on back-bone links via packet header analysis. *Campus-wide Inf. Syst.* 25, 5 (2008), 342–358.

[95] Arnold Johnson, Kelley Dempsey, Ron Ross, Sarbari Gupta, and Dennis Bailey. 2011. Guide for security-focused configuration management of information systems. Technical Report SP 800-128, U.S. Dept of Commerce, National Institute of Standards and Technology, Gaithersburg, MD.

[96] Joint Chiefs of Staff. 2014. Joint intelligence preparation of the operational environment. Technical Report JP 2-01.3, U.S. Dept. of Defense, Washington, DC.

[97] Joint Chiefs of Staff. 2012. Information operations. Technical Report JP 3-13, U.S. Dept. of Defense, Washington, DC.

[98] Joint Chiefs of Staff. 2013. Joint targeting. Technical Report JP 3-60, U.S. Dept. of Defense, Washington, DC.

[99] Vaibhavi Kalgutkar, Ratinder Kaur, Hugo Gonzalez, Natalia Stakhanova, and Alina Matyukhina. 2019. Code authorship attribution: Methods and challenges. *ACM Comput. Surv.* 52, 1 (2019), 3:1–3:36. DOI: http://doi.acm.org/10.1145/3292577

[100] Jonathan Katz. 2016. Call for papers: Hot topics in the science of security (HoTSoS). Retrieved from http://cps-vo.org/group/hotsos/cfp.

[101] Karen Kent and Murugiah Souppaya. 2006. Guide to computer security log management. Technical Report SP 800-92, U.S. Dept of Commerce, National Institute of Standards and Technology, Gaithersburg, MD.

[102] Karen Kent, Suzanne Chevalier, Tim Grance, and Hung Dang. 2006. Guide to integrating forensic techniques into incident response. Technical Report SP 800-86, U.S. Dept. of Commerce, National Institute of Standards and Technology, Gaithersburg, MD.

[103] Suleman Khan, Abdullah Gani, Ainuddin Wahid Abdul Wahab, Mustapha Aminu Bagiwa, Muhammad Shiraz, Samee U. Khan, Rajkumar Buyya, and Albert Y. Zomaya. 2016. Cloud log forensics: Foundations, state of the art, and future directions. *ACM Comput. Surv.* 49, 1 (2016), 7:1–7:42.

[104] Klaus-Peter Kossakowski, William R. Wilson, Julia H. Allen, Cecilia Albert, Cory Cohen, Gary Ford, Barbara Fraser, Eric Hayes, John Kochmar, and Suresh Konda. 1999. Responding to intrusions. Technical Report CMU/SEI-99-SIM-006, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA.

[105] Kat Krol, Jonathan M. Spring, Simon Parkin, and M. Angela Sasse. 2016. Towards robust experimental design for user studies in security and privacy. In *Learning from Authoritative Security Experiment Results (LASER'16)*. IEEE, 21–31.

[106] Dirk Kuhlmann, Liqun Chen, and Christopher J. Mitchell. 2016. Trust and legitimacy in security standardization—A new management issue? In *Interoperability for Enterprise Systems and Applications (I-ESA'16)*. ISTE Publications, Guimaraes, Portugal.

[107] Ruggero Donida Labati, Angelo Genovese, Enrique Muñoz, Vincenzo Piuri, Fabio Scotti, and Gianluca Sforza. 2016. Biometric recognition in automated border control: A survey. *ACM Comput. Surv.* 49, 2 (2016), 24:1–24:39.

[108] Aron Laszka, Mark Felegyhazi, and Levente Buttyan. 2014. A survey of interdependent information security games. *ACM Comput. Surv.* 47, 2 (2014), 23:1–23:38.

[109] Stefan Laube and Rainer Böhme. 2017. Strategic aspects of cyber risk information sharing. *ACM Comput. Surv.* 50, 5 (2017), 77:1–77:36. DOI: http://doi.acm.org/10.1145/3124398

[110] M. Leech. 2003. Chinese Lottery Cryptanalysis Revisited: The Internet as a Codebreaking Tool. RFC 3607 (Informational). https://tools.ietf.org/html/rfc3607.

[111] Ryan Leigland and Axel W. Krings. 2004. A formalization of digital forensics. *Int. J. Dig. Evid.* 3, 2 (2004), 1–32.

[112] James A. Lewis. 2008. *Holistic Approaches to Cybersecurity to Enable Network Centric Operations*, Vol. 1. Center for Strategic and International Studies.

[113] Tao Li, Ning Xie, Chunqiu Zeng, Wubai Zhou, Li Zheng, Yexi Jiang, Yimin Yang, Hsin-Yu Ha, Wei Xue, Yue Huang, Shu-Ching Chen, Jainendra Navlakha, and S. S. Iyengar. 2017. Data-driven techniques in disaster information management. *ACM Comput. Surv.* 50, 1 (2017), 1:1–1:45.

[114] Tao Li, Chunqiu Zeng, Yexi Jiang, Wubai Zhou, Liang Tang, Zheng Liu, and Yue Huang. 2017. Data-driven techniques in computing system management. *ACM Comput. Surv.* 50, 3 (2017), 45:1–45:43.

[115] Huan Liu and Hiroshi Motoda. 1998. *Feature Selection for Knowledge Discovery and Data Mining*. Springer Science & Business Media, New York.

[116] Ming Liu, Zhi Xue, Xianghua Xu, Changmin Zhong, and Jinjun Chen. 2018. Host-based intrusion detection system with system calls: Review and future trends. *ACM Comput. Surv.* 51, 5 (2018), 98:1–98:36. DOI: http://doi.acm.org/10.1145/3214304

[117] Jason T. Luttgens, Matthew Pepe, and Kevin Mandia. 2014. *Incident Response & Computer Forensics* (3rd ed.). McGraw-Hill Education Group.

[118] G. B. Magklaras and S. M. Furnell. 2001. Insider threat prediction tool: Evaluating the probability of it misuse. *Comput. Sec.* 21, 1 (2001), 62–73.

[119] Mandiant. 2013. APT1: Exposing one of China's cyber espionage units. Technical Report. FireEye, Inc.

[120] Stuart McClure, Joel Scambray, and George Kurtz. 2005. *Hacking Exposed: Network Security Secrets & Solutions* (5th ed.). McGraw-Hill Osborne.

[121] Guozhu Meng, Yang Liu, Jie Zhang, Alexander Pokluda, and Raouf Boutaba. 2015. Collaborative security: A survey and taxonomy. *ACM Comput. Surv.* 48, 1 (2015), 1:1–1:42.

[122] Leigh B. Metcalf and Jonathan M. Spring. 2015. Blacklist ecosystem analysis: Spanning Jan 2012 to Jun 2014. In *the 2nd ACM Workshop on Information Sharing and Collaborative Security*. 13–22.

[123] Aleksandar Milenkoski, Marco Vieira, Samuel Kounev, Alberto Avritzer, and Bryan D. Payne. 2015. Evaluating computer intrusion detection systems: A survey of common practices. *ACM Comput. Surv.* 48, 1 (2015), 12:1–12:41.

[124] George A. Miller. 1956. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychol. Rev.* 63, 2 (1956), 81.

[125] Robin Milner. 1989. *Communication and Concurrency*. Prentice Hall, New York.

[126] MITRE Corporation. 2010. Science of cyber-security. Technical Report JSR-10-102, JASON Office, McLean, VA.

[127] Sarandis Mitropoulos, Dimitrios Patsos, and Christos Douligeris. 2006. On incident handling and response: A state-of-the-art approach. *Comput. Sec.* 25, 5 (2006), 351–370.

[128] K. Moriarty. 2010. Real-time Inter-network Defense (RID). RFC 6045 (Informational). Retrieved from https://www.rfc-editor.org/rfc/rfc6045.txt.

[129] K. Moriarty. 2012. Real-time Inter-network Defense (RID). RFC 6545 (Proposed Standard). https://tools.ietf.org/html/rfc6545.

[130] K. Moriarty and B. Trammell. 2010. Transport of Real-time Inter-network Defense (RID) Messages. RFC 6046 (Informational). Retrieved from https://www.rfc-editor.org/rfc/rfc6046.txt.

[131] David A. Mundie and Robin Ruefle. 2012. Building an incident management body of knowledge. In *Availability, Reliability and Security (ARES'12)*. IEEE, 507–513.

[132] David A. Mundie, Robin Ruefle, Audrey J. Dorofee, Samuel J. Perl, John McCloud, and Matthew Collins. 2014. An incident management ontology. In *Semantic Technology for Intelligence, Defense, and Security*. C4I, Fairfax, VA, 62–71.

[133] Paul Nightingale. 2009. Tacit knowledge and engineering design. In *Philosophy of Technology and Engineering Sciences*, Handbook of the Philosophy of Science. North-Holland, Amsterdam, 351–374.

[134] Peter W. O'Hearn. 2007. Resources, concurrency, and local reasoning. *Theor. Comput. Sci.* 375, 1 (2007), 271–307.

[135] Steven Oksala, Anthony Rutkowski, Michael Spring, and Jon O'Donnell. 1996. The structure of IT standardization. *StandardView* 4, 1 (1996), 9–22.

[136] Tore Larsen Orderløkken. 2005. Security incident handling and reporting—A study of the difference between theory and practice. Master's thesis. Gjøvik University College, Norway.

[137] Marcos Osorno, Thomas Millar, and Danielle Rager. 2011. Coordinated cybersecurity incident handling: Roles, processes, and coordination networks for crosscutting incidents. Technical Report, Johns Hopkins University, Applied Physics Laboratory, Laurel, MD.

[138] Gary Palmer. 2001. A road map for digital forensic research. In *First Digital Forensic Research Workshop* 27–30.

[139] Michael J. Palmiotto (Ed.). 1988. *Crime Pattern Analysis: An Investigative Tool* (2nd ed.). Pilgrimage, 59–69.

[140] Min-Seok Pang and Huseyin Tanriverdi. 2017. Security breaches in the U.S. federal government. In *Workshop on the Economics of Information Security.*

[141] Michael Pearce, Sherali Zeadally, and Ray Hunt. 2013. Virtualization: Issues, security threats, and solutions. *ACM Comput. Surv.* 45, 2 (2013), 17:1–17:39.

[142] Marcus Pendleton, Richard Garcia-Lebron, Jin-Hee Cho, and Shouhuai Xu. 2016. A survey on systems security metrics. *ACM Comput. Surv.* 49, 4 (2016), 62:1–62:35.

[143] Tao Peng, Christopher Leckie, and Kotagiri Ramamohanarao. 2007. Survey of network-based defense mechanisms countering the dos and ddos problems. *ACM Comput. Surv.* 39, 1 (2007).

[144] Mark Pollitt. 2008. Applying traditional forensic taxonomy to digital forensics. In *Advances in Digital Forensics*, Indrajit Ray and Sujeet Shenoi (Eds.). 17–26.

[145] David Pym. 2018. The origins of cyberspace. In *Oxford Handbook of Cyber Security.* OUP.

[146] Ali Ahmadian Ramaki, Abbas Rasoolzadegan, and Abbas Ghaemi Bafghi. 2018. A systematic mapping study on intrusion alert analysis in intrusion detection systems. *ACM Comput. Surv.* 51, 3 (2018), 55:1–55:41. DOI: http://doi.acm.org/10.1145/3184898

[147] Mark Reith, Clint Carr, and Gregg Gunsch. 2002. An examination of digital forensic models. *Int. J. Dig. Evid.* 1, 3 (2002), 1–12.

[148] Richard L. Rollason-Reese. 2003. Incident handling: An orderly response to unexpected events. In *the 31st ACM SIGUCCS Fall Conference.* 97–102.

[149] Ron Ross, Gary Stoneburner, Richard Graubart, Kelley Dempsey, Esten Porter, Bennett Hodge, Karen Quigg, Christian Enloe, Kevin Stine, Jennifer Fabius, Daniel Faigin, Arnold Johnson, Lisa Kaiser, Pam Miller, Sandra Miravalle, and Victoria Pillitteri. 2013. Security and privacy controls for federal information systems and organizations. Technical Report SP 800-53r4, U.S. Dept. of Commerce, National Institute of Standards and Technology, Gaithersburg, MD.

[150] Christian Rossow, Christian J. Dietrich, Chris Grier, Christian Kreibich, Vern Paxson, Norbert Pohlmann, Herbert Bos, and Maarten Van Steen. 2012. Prudent practices for designing malware experiments: Status quo and outlook. In *the IEEE Symposium on Security and Privacy (S&P'12)*. 65–79.

[151] Robert Rowlingson. 2004. A ten-step process for forensic readiness. *Int. J. Dig. Evid.* 2, 3 (2004), 1–28.

[152] Arpan Roy, Santonu Sarkar, Rajeshwari Ganesan, and Geetika Goel. 2015. Secure the cloud: From the perspective of a service-oriented organization. *ACM Comput. Surv.* 47, 3 (2015), 41:1–41:30.

[153] Karen Scarfone and Peter Mell. 2007. Guide to intrusion detection and prevention systems (IDPS). Technical Report SP 800-94, U.S. Dept. of Commerce, National Institute of Standards and Technology, Gaithersburg, MD.

[154] Karen Scarfone, Murugiah Souppaya, Amanda Cody, and Angela Orebaugh. 2008. Technical guide to information security testing and assessment. Technical Report SP 800-115, U.S. Dept of Commerce, National Institute of Standards and Technology, Gaithersburg, MD.

[155] Piya Shedden, Atif Ahmad, and Anthonie B. Ruighaver. 2011. Informal learning in security incident response teams. In *the Australasian Conference on Information Systems*.

[156] R. Shirey. 2007. Internet Security Glossary, Version 2. RFC 4949 (Informational). https://tools.ietf.org/html/rfc4949.

[157] Herbert A. Simon. 1996. *The Sciences of the Artificial* (3rd ed.). The MIT Press, Cambridge, MA.

[158] Rebecca Slayton and Brian Clarke. 2020. Trusting infrastructure: The emergence of computer security incident response, 1989–2005. *Technol. Cult.* 61, 1 (2020), 173–206.

[159] Danny Smith and Moira West-Brown. 1996. Incident handling—Experience through role-playing. In *FIRST Conference and Workshop on Computer Security Incident Handling and Response.* Retrieved from https://www.first.org/conference/1996/tutorials.html#tuta.

[160] Muragiah Souppaya and Karen Scarfone. 2013. Guide to malware incident prevention and handling for desktops and laptops. Technical Report SP 800-83r1, U.S. Dept. of Commerce, National Institute of Standards and Technology, Gaithersburg, MD.

[161] Jonathan M. Spring. 2013. Modeling malicious domain name take-down dynamics: Why eCrime pays. In *eCrime Researchers Summit (eCRS'13).* IEEE.

[162] Jonathan M. Spring and Eric Hatleback. 2016. Thinking about intrusion kill chains as mechanisms. *J. Cybersec.* 3 (3), 185–197.

[163] Jonathan M. Spring and Phyllis Illari. 2018. Building general knowledge of mechanisms in information security. *Philos. Technol.* DOI: 10.1007/s13347-018-0329-z

[164] Jonathan M. Spring, Sarah Kern, and Alec Summers. 2015. Global adversarial capability modeling. In *APWG Symposium on Electronic Crime Research (eCrime'15).* IEEE.

[165] Jonathan M. Spring, Tyler Moore, and David Pym. 2017. Practicing a science of security: A philosophy of science perspective. In *New Security Paradigms Workshop*.

[166] Michael B. Spring. 2011. What have we learned about standards and standardization? *Homo Oeconom.* 27, 4 (2011), 501–517.

[167] Alex Stamos. 2010. "Aurora" response recommendations. Technical Report, iSec Partners.

[168] Clifford Stoll. 1988. Stalking the wily hacker. *Commun. ACM* 31, 5 (1988), 484–497.

[169] Blake E. Strom, Joseph A. Battaglia, Michael S. Kemmerer, William Kupersanin, Douglas P. Miller, Craig Wampler, Sean M. Whitley, and Ross D. Wolf. 2017. Finding cyber threats with ATT&CK-based analytics. Technical Report MTR17-0202, MITRE Corporation, Annapolis Junction, MD.

[170] Blake E. Strom, Andy Applebaum, Doug P. Miller, Kathryn C. Nickels, Adam G. Pennington, and Cody B. Thomas. 2018. MITRE ATT&CK: Design and philosophy. Technical Report MP18-0360, MITRE Corporation, McLean, VA.

[171] Sathya Chandran Sundaramurthy, John McHugh, Xinming Simon Ou, S. Raj Rajagopalan, and Michael Wesch. 2014. An anthropological approach to studying CSIRTs. *IEEE Sec. Priv.* 5 (2014), 52–60.

[172] SWGDE. 2009. Position on the National Research Council report to Congress strengthening forensic science in the United States: A path forward. Retrieved from https://www.swgde.org/documents/Current%20Documents/SWGDE%20Position%20on%20the%20NAS%20Report.

[173] T. Takahashi, K. Landfield, and Y. Kadobayashi. 2014. An Incident Object Description Exchange Format (IODEF) Extension for Structured Cybersecurity Information. RFC 7203 (Proposed Standard). https://tools.ietf.org/html/rfc7203.

[174] Jun Tang, Yong Cui, Qi Li, Kui Ren, Jiangchuan Liu, and Rajkumar Buyya. 2016. Ensuring security and privacy preservation for cloud data services. *ACM Comput. Surv.* 49, 1 (2016), 13:1–13:39.

[175] John A. Tirpak. 2000. Find, fix, track, target, engage, assess: F2T2EA is shorthand for the operational goal the air force will pursue into the 21st century. *Air Force Mag.* 83, 7 (2000), 24–29.

[176] B. Trammell. 2012. Transport of Real-time Inter-network Defense (RID) Messages over HTTP/TLS. RFC 6546 (Proposed Standard). https://tools.ietf.org/html/rfc6546.

[177] B. Trammell. 2012. Guidelines and Template for Defining Extensions to the Incident Object Description Exchange Format (IODEF). RFC 6684 (Informational). https://tools.ietf.org/html/rfc6684.

[178] U.S. Dept of Commerce. 2017. NIST mission, vision, core competencies, and core values. Retrieved from https://www.nist.gov/about-nist/our-organization/mission-vision-values.

[179] Aleksandar Valjarevic and Hein S. Venter. 2012. Harmonised digital forensic investigation process model. In *Information Security for South Africa (ISSA'12)*. IEEE, 1–10.

[180] Aleksandar Valjarevic and Heini S. Venter. 2012. Analyses of the state-of-the-art digital forensic investigation process models. In *Southern Africa Telecommunication Networks and Applications Conference (SATNAC)*, Stefan Scriba (Ed.).

[181] Wim Van Eck. 1985. Electromagnetic radiation from video display units: An eavesdropping risk? *Comput. Sec.* 4, 4 (1985), 269–286.

[182] Michael Vangelos. 2010. Incident response: Managing. In *Encyclopedia of Information Assurance*. Auerbach Publications, 1442–1449.

[183] Verizon. 2015. 2015 data breach investigations report (DBIR). Technical Report. Retrieved from http://www.verizonenterprise.com/DBIR/2015/.

[184] Verizon. 2016. 2016 data breach investigations report (DBIR). Technical Report. Retrieved from http://www.verizonenterprise.com/verizon-insights-lab/dbir/2016/.

[185] Ju An Wang and Minzhe Guo. 2009. OVM: An ontology for vulnerability management. In *Workshop on Cyber Security and Information Intelligence Research*. ACM.

[186] Rodrigo Werlinger, David Botta, and Konstantin Beznosov. 2007. Detecting, analyzing and responding to security incidents: A qualitative analysis. In *the 3rd Symposium on Usable Privacy and Security*. 149–150.

[187] Janet Williams. 2012. Good practice guide for digital evidence. Technical Report, Association of Chief Police Officers, London, U.K.

[188] Robert Willison and Mikko Siponen. 2009. Overcoming the insider: Reducing employee computer crime through situational crime prevention. *Commun. ACM* 52, 9 (2009), 133–137.

[189] Yanfang Ye, Tao Li, Donald Adjeroh, and S. Sitharama Iyengar. 2017. A survey on malware detection using data mining techniques. *ACM Comput. Surv.* 50, 3 (2017), 41:1–41:40.

[190] Yves Younan, Wouter Joosen, and Frank Piessens. 2012. Runtime countermeasures for code injection attacks against C and C++ programs. *ACM Comput. Surv.* 44, 3 (2012), 17:1–17:28.

Received September 2019; revised June 2020; accepted October 2020