# Rethinking Information Sharing for Threat Intelligence

[Position Paper]

Aziz Mohaisen
University of Central Florida
mohaisen@ucf.edu

Omar Al-Ibrahim
University at Buffalo
omaralib@buffalo.edu

Charles Kamhoua
Army Research Laboratory
charles.a.kamhoua.civ@mail.mil

Kevin Kwiat
Air force Research Lab
kevin.kwiat@us.af.mil

Laurent Njilla
Air force Research Lab
laurent.njilla@us.af.mil

## ABSTRACT

In the past decade, the information security and threat landscape has grown significantly making it difficult for a single defender to defend against all attacks at the same time. This called for introducing information sharing, a paradigm in which threat indicators are shared in a community of trust to facilitate defenses. Standards for representation, exchange, and consumption of indicators are proposed in the literature, although various issues are undermined. In this paper, we take the position of rethinking information sharing for actionable intelligence, by highlighting various issues that deserve further exploration. We argue that information sharing can benefit from well-defined use models, threat models, well-understood risk by measurement and robust scoring, well-understood and preserved privacy and quality of indicators and robust mechanism to avoid free riding behavior of selfish agents. We call for using the differential nature of data and community structures for optimizing sharing designs and structures.

## CCS CONCEPTS

• **Security and privacy → Vulnerability management**;

## KEYWORDS

Threat intelligence, information sharing, standards, privacy

**ACM Reference format:**
Aziz Mohaisen, Omar Al-Ibrahim, Charles Kamhoua, Kevin Kwiat, and Laurent Njilla. 2017. Rethinking Information Sharing for Threat Intelligence. In *Proceedings of HotWeb'17, San Jose / Silicon Valley, CA, USA, October 14, 2017,* 7 pages.
https://doi.org/10.1145/3132465.3132468

## 1 INTRODUCTION

New information and communication technology platforms, such as cloud and mobile computing, social networks, and the Internet of Things (IoT), reshaped the security landscape to become more sophisticated. Unorganized forms of vandalism have become a diverse ecosystem of cybercrime, where providers and consumers come together to achieve various end-goals and utilities [18, 19]. The persistence, complexity, size, and capabilities of today's adversaries are unbounded, and their threat does not only affect individuals or organizations, but also nations as a whole: according to a recent study [14], direct and indirect costs due to security breaches have costed the global economy about $491 billion in 2014 alone.

Visibility into behaviors and capabilities of adversaries to form detection signatures is an essential first step towards containing and defending against them, and ultimately thwarting their harms. On the other hand, with the unprecedented complexity and size of the threat ecosystem, no single defender can defend against all attacks all the time. Even when facing attacks, defenders need to have the right skills to recognize them before performing defense efforts. With the skill gap on the rise, visibility into attacks and malicious actors becomes a challenge. Thus, a coordinated solution based upon the collective knowledge of multiple defenders is required. In such a solution, multiple stakeholders share information about security incidents observed and collected from their security operations, with the hope that such information would be useful to other stakeholders in improving their security posture.

### 1.1 Information Sharing

Information sharing has emerged as a plausible solution to addressing the aforementioned problems. Threat information sharing is utilized for efficiently and effectively defending against emerging threats. One even went as far as to say that "threat intelligence sharing is the only way to combat the growing skills gap" [10]. In practice, information sharing is used to communicate operational security experience between participants in a sharing system with the hope that sharing would 1) enable participants to defend their systems against ongoing attacks, and 2) improve their defense posture by proactively addressing possible attacks.

Information sharing is not a theoretical idea, and there is a lot of work on defining tools for representing information, or mechanisms for exchanging such information between information sharing participants in sharing communities. Information sharing also has been embraced by various communities, and leaders in such community have created their own sharing exchange points, where participants could deliver and retrieve the shared raw data and annotated data (intelligence) from other participants using standard application program interfaces (APIs): for example, Facebook


*HotWeb'17, October 14, 2017, San Jose / Silicon Valley, CA, USA*

ACM ISBN 978-1-4503-5527-8/17/10…$15.00
https://doi.org/10.1145/3132465.3132468

has created ThreatExchange [4], and Verisign has created the IntelGraph, among others. Such initiatives are not limited to the private sector: the public sector initiated information sharing in the US Department of Homeland Security's (DHS) Critical Infrastructure Cyber Information Sharing and Collaboration (CISCP) program [1], which aims to facilitate sharing of threat indicators between the private and public sectors and vice versa.

### 1.2 Risks of Sharing

The risk of not sharing information is clear, which can be seen in more and more security breaches using the same attack vectors and capabilities, and re-examining the same vulnerabilities. Despite the various benefits of information sharing for security, even within a limited community of participants, shared information without proper restrictions, however, may leak a significant amount of information about the participants and their operation context. For example, information shared for the good and security of the participating community can be also used by the adversary to learn about vulnerabilities of those participants. Also, the same vulnerabilities can be used to test their applicability on other systems: with the lag in patching vulnerabilities, the adversary will be able to utilize such information for attacking other unpatched systems.

This risk of sharing can be mitigated in information sharing by limiting the sharing community to highly trusted participants, and informing potentially subjected participants with the risk ahead of time, and certainly before sharing such information with a broader community. However, limiting sharing to a highly trusted community of participants in general security applications, while reducing risk of information exposure to unintended parties—including adversaries—may have an equally damaging consequence: the security of today's systems with the presence of multiple stakeholders exhibit the end-to-end principle of design characterized by fate-sharing. For example, an unpatched system under the control of uninformed player simply because that player is not trusted (enough) can be used to attack other systems under the control of the highly trusted participants. Classification of participants' risk in sharing communities is provided in the literature, despite the lack clarity on how such risk should be assessed [2].

Another scenario of sharing where negative consequences may arise is privacy of individuals, and how sharing may affect civil liberties [6]. The sharing of public data that does not in of and itself identify individuals would serve the goal of information sharing without any side effects on privacy. However, it is believed that privacy does not often go along well with security: to be able to attribute attacks and perform various security analyses, context information should be present along with the threat indicators for further inferences that would serve security [20]. For example, along with an end point (i.e., hostname, or IP address) an incident indicator typically shared would include e-mail addresses, URLs, etc., from which intelligence is gathered, and security risk is assessed.

Privacy risks due to sharing are arguably mitigated by a minimalistic approach, where only a limited amount of data is *collected and shared* [2]. However, whether such a minimalistic approach is being implemented in today's sharing paradigms or not is unclear. Furthermore, such an approach goes against security utilities. We conjecture (with confidence based on various plausible applications) that the additional context information of the threat information shared is often times as important as the indicators themselves. To this end, a new approach to thinking privacy is required beyond simple minimization. Such a technique could perhaps utilize techniques for safeguarding all necessary information to improve the security posture of a defender, while ensuring the privacy of users, and confidentiality of shared information.

Related to the community of trust problem above and perceived risk of over sharing (whether it is for security or privacy) is the problem of "free-riding". As a result of the perceived risk of sharing, some actors might be actually joining communities of sharing, although not sharing sufficient information to the community that others can benefit from it [7, 8]. When a community member joins and shares information there is always the risk of the shared information (e.g., about a vulnerability) being leaked to the public (or even worse, to the adversary), resulting in both monetary and reputation loses. Such scenario leads to that some actors might not truthfully share information due to their own self-interests. While recent works have been focused on addressing problem in a theoretical framework [16, 17], assuming the level of participation as an indicator of contribution in information sharing, there is no work that extends beyond that to account for quality of indicators. For example, an actor that contributes stale indicators, indicators that are not timely to be utilized operationally, while not considered a free-rider in the typical sense, is not contributing sufficiently and meaningfully to the missing of information sharing.

**Objectives.** Believing in their beneficial aspects, the goal of this work is to shed light on various issues associated with information sharing, including understanding community structures, use and adversary models, privacy issues and quality of indicators for detecting free riding in information sharing for actionable threat intelligence. With standard sharing formats being widely advocated as the next step towards effective sharing, we identify the need for understanding privacy and risk. To understand this risk in context, we identify a plausible sharing scenario for which we define the adversary models associated with both internal or external adversaries. We introduce to the analysis the various sharing paradigms under them. By identifying the need for security, we advocate an approach that combines various aspects of design techniques that exploit the differential nature of data and community structure. Finally, we identify quality of indicators as an important direction, suggest various directions to assessing quality, and call the research community to further the suggested directions.

**Organization.** The rest of this paper is organized as follows. In section 2, we provide our broad vision of various directions for rethinking sharing towards actionable threat intelligence. In section 3 we ellaborate on one of directions, namely privacy. In section 4 we elaborate on another issue, namely, quality of indicators. Concluding remarks are in section 5.

## 2 RETHINKING SHARING

Realizing actionable intelligence by striking a balance between utility of the information sharing systems and other requirements, including privacy, security, and complexity of the sharing system, is a non-trivial task. In the following, we offer to rethink sharing by touching on various fundamental issues and building blocks in

typical sharing systems. We identify the following issues as rich areas that require further research and exploration, and offer various directions associated with each of those issues in the subsequent sections. We offer to understand use models (§2.1), sharing communities (§2.2), adversaries in sharing paradigms, including both outsider and insider adversaries (§2.3), quantifying understanding privacy in information sharing, towards measurement.

## 2.1 Defining the Use Model

Information sharing is inseparable from its use model and scenario. Thus, understanding the various technical details of the use model of information sharing tools and paradigms is essential to understanding various issues, including security, privacy, and functional issues. We offer to touch on various scenarios of use and issues associated with in the following.

We classify the use models of information sharing for threat intelligence into various types based on various classification criteria, as follows. 1) Structure: Based on the *structure* and format of the shared information, we classify information sharing tools into structured (standard) and unstructured sharing models. 2) Centralization: Based on whether a centralized sharing entity (repository) exists or not, we classify such models into centralized and decentralized systems. 3) Scope and function: Information sharing tools can be also classified based on their scope and function. While it is difficult to enumerate such scopes and functions for unstructured sharing systems, structured systems that use standards are classified into enumeration, scoring, languages, and transport mechanisms. More details are provided in §3.

**Unstructured Sharing.** The end goal of information sharing is to realize a secure cyberspace by exchanging operational security experience across multiple players in a sharing community. Whether data used in the information sharing paradigm is structured or not is irrelevant to the main goal of information sharing. Traditionally, threat information concerning incidents has been collected and shared as unstructured data, and exchanged via generic communication tools and services, including electronic mail, or file transfer services. Today, and despite the rise of structure via standardized formats and sharing schemas, proprietary formats are widely used in by vendors in security market, making interoperation between structures hard to achieve. While it is easier to understand structured schemas, where various attributes are indicated, understanding the privacy of sharing when using unstructured formats is not possible. To this end, in the rest of this work we focus on structured sharing format, although we believe that unstructured sharing also may have various privacy risks that should be studied and addressed based on actual assessments.

**Sharing Using Standard Formats.** For efficient use of shared information in an automated manner, it is desirable to share information in a standard and structured format. For that, there has been a lot of work in the literature on understanding use scenarios, and developing the relevant schemas of structured formats for information sharing. By understanding the type of data in such information sharing formats, it would not only be possible to understand the capabilities embodied in the various sharing formats, but also to understand the privacy risks in the abstract, and possibly develop technical solutions to address it. Examples of such sharing paradigms include CVE, CCE, CWE, CyBox, etc. More on those schemas is in §3.1.

## 2.2 Sharing Communities

Sharing is defined in "communities of trust", which are the structure in which threat information is shared to reach a common goal of strengthening the security posture of various participants in the community. Sharing today is defined based on the nature of the participants (whether they are public or private sector participants) into private-private, public-private, and public-public. An example of the private-private sector information sharing communities include participants in the likes of ThreatExchange, or IntelGraph, while an example of the public-private partnerships include DHS's CISCP [1].

On the one hand, various of those communities are vetted carefully to ensure that the information being shared between the various participants in the sharing system is safeguarded and not used to attack any participant in the system. On the other hand, circumstantial evidence (or even conclusive evidence [15]) has shown that information being shared in the sharing system could potentially be used as an attack vector against another participant in the system. Understanding the make up of the sharing community is perhaps a first step to account for such risk.

**Redefining communities.** Redefining communities structure by relaxing assumptions of "trust" in a way that would allow for a greater participation of players in a sharing system results, thus potentially resulting for improved defenses and security awareness by a larger number of participants, would potentially result in a higher risk of sharing. Such risk is not only increasing the attack surface, but potentially disincentivizing major community members from meaningful and sharing of quality information resulting in actionable intelligence. Understanding how relaxing the definition of communities would affect both utility of sharing and the risk is to be explored further in light of actual and measurable risk.

**Privacy-based community definition.** So far, communities have been defined for their trustworthiness with respect to their risk awareness, or for utilizing the various tools and paradigms of information sharing, but not understood w.r.t. privacy. Thus, we believe it is a worthwhile to incorporate privacy as a metric (along with other metrics of risk or in isolation) as a criteria for defining communities. Furthermore, technical solutions that take into account a clear definition of privacy-awareness and its presence (or lack) in a certain community (or players in a community) could be further optimized to suite the underlying assumptions of such community.

## 2.3 Understanding Adversaries

Security and privacy of communication and computation protocols are often analyzed under various settings of adversaries. Adversaries are characterized by capabilities under which security and privacy definitions are formalized, and security and privacy guarantees (in light of a formally defined advantage of the adversary) are outlined. With the complexity and involved nature of information sharing paradigms, and the end-goal that they try to achieve, we argue that both insider and outsider (external) adversaries are relevant to studying the information sharing in the field. In the following, we elaborate on both forms of threat, and open directions to address in order to realize a formally-backed exchange.

**External adversary.** Such adversaries are defined broadly as adversaries who are not part of the system or protocol being analyzed, and they include various forms of actors, ranging simply from a passive eavesdropper [3, 5] or honest-but-curious to the more advanced active adversary–an adversary that could potentially interfere with communication or manipulate computations in order to affect the security of the system, or breach the privacy of a participant. This adversary can be a single malicious actor, or multiple of them. The main qualifier of this adversary, however, is that it is not included in the set of participants of the system.

Instances of such adversary include simple observers on the communication channel between participants in the information sharing system, with their risks being mitigated by the various in-place cryptographic techniques. Another example of the observer could be a publicly shared infrastructure, like cloud, where the cloud provider may have a great incentive not to act maliciously, but would be interested in knowing some details about the information being shared and hosted in the cloud. While auditing and strict policies are one direction to tame this adversary, relinquishing trust and enforcing a stronger form of audit—perhaps by utilizing cryptographic approaches, is yet another method. We elaborate on such methods in the subsequent sections. The aforementioned example of cloud could be also viewed as a totally untrusted, and potentially malicious, thus being an instance of the malicious adversary. Such state of being malicious could be a property of the cloud itself; i.e., the cloud provider is untrusted, or due to other externalities, e.g., the cloud is being compromised by an outsider through, for example, a malware campaign [11]. The way that such adversary is realized is irrelevant to understanding the privacy of the various sharing paradigms, although the capabilities of such adversary are.

**Insider adversary.** Motivated by the various risks that potentially could be the result of misuse of the information shared an information sharing system [15], another adversary model that needs to be formalized is the insider adversary. Whereas typical threats in various systems include the external adversary highlighted above, the nature of information sharing systems highlight that insider adversaries are real risks. Such adversaries could be in multiple forms, and stem out of various system and operation realities. For example, such adversary could be another participant in the information sharing system acting maliciously to reach a certain objective, or an individual acting on behalf of a participant in the system. Understanding how information sharing is prone to such class of adversaries is necessary to enable sharing. Furthermore, such adversary could perhaps be studied well under other notions of risk associated with information sharing and definition of communities of trust, their risk and privacy awareness.

## 2.4 Evidence-based Analysis

One may argue the problem at hand is not any different from any other privacy problem due to data exposure, thus thinking of the privacy issues in information sharing for threat intelligence in the abstract is meaningful and the way this problem should be addressed given the large number of use scenarios.

We argue that while thinking of this problem in the abstract is worthwhile, also approaching the problem with technical solutions that stems from the actual size and shape of privacy exposure in the various information sharing paradigms is important. A first step towards understanding the actual size and shape of exposure is facilitated by an actual quantification of exposure in real data. However, one cannot quantify what he cannot measure, thus measuring data exposure in the various sharing paradigms, under the various settings of threat models or in isolation, is necessary and important for understanding the problem in practice. In particular, measurements would give abstract studies context that highlight actual findings related to indicators, privacy, and risk.

Measuring privacy leakage in the various paradigms of sharing and under various models is not an easy task. We argue that privacy cannot be understood in the abstract, and without a clear context of sharing [13]. Even worse, what constitute a privacy concern to one individual might not be of value to another individual. Thus, a first step to measuring privacy leakage in information sharing is to formalize what we mean by privacy, what are the private attributes that should be treated with care and hidden from adversaries and other (potentially honest-but-curious) participants, and how sensitive (with respect to their privacy value) alone or when associated with other data about the subject.

## 2.5 Quality and Privacy

Quality of the indicators and privacy are at odds: in order to provide the highest accuracy in security operations, access to raw and highlight annotated indicators that can be of use for actionable intelligence is necessary. On the other hand, having such raw indicators without any sanitization or masking of any of their contents could potentially leak the privacy of entities associated, or reveal sensitive information about the operation context where they are collected, directly or indirectly. To this end, another direction to pursue is by answering the following question: How much quality of indicators should be given up to satisfy various privacy notions and guarantees.

This question is not easy to answer: there are various competing and varying notions of privacy, and systematically and formally analyzing and modeling how they are met (or violated) at various levels of exposure of indicators. Before even approaching this question, it would be necessary to formalize metrics for evaluating the quality of the indicators.

# 3 UNDERSTANDING RISK IN SHARING

There is a clear risk of sharing, whether it is privacy or security. Understanding such risk is the first step towards providing practical solutions to the various aspects of risk. In the following, we elaborate on a road-map for understanding risk in information sharing, mainly emphasizing privacy risks. In §3.1 we review the various sharing schemes. In §3.2 we highlight risks of information sharing through various measurements and examples from *anonymized* sharing datasets. In §3.3 we argue for a privacy leakage assessment design that takes into account the various issues raised on the risk of the sharing paradigms. In §3.4 we advocate architectural design accounting privacy and community structure as a design principle.

## 3.1 An Overview of Sharing Standards

As noted previously, there are various standards for information sharing that are used by government and industry to automate and structure the exchange of information within an organization and

between autonomous systems and organizations. We can classify these sharing standards into four main categories:

**Enumerations.** Standardized enumerations of platforms, configurations, software weaknesses, and attacks. Examples include Common Configuration Enumeration (CCE), Common Weakness Enumeration (CWE), and Common Vulnerabilities and Exposures (CVE).

**Scoring systems.** Standards to assess the severity of computer system-related issues and assigning scores to each one, allowing responders to prioritize remediation tasks. Common standards that fit this category include Common Vulnerability Scoring System (CVSS) and Common Weakness Scoring System (CWSS).

**Languages.** Those standards are intended for encoding high-fidelity information about systems in a manner that facilitates parsing this information in software tools and converting them to human-readable formats. This includes formats like the Malware Attribute Enumeration and Characterization (MAEC), Open Vulnerability and Assessment Language (OVAL), Incident Object Description Exchange Format (IODEF), Extensible Configuration Checklist Description Format (XCCDF), and Structured Threat Information Exchange (STIX).

**Transport.** Those standards represent Inter-network communication formats to facilitate exchange of information between hosts. Standards such as Real-time Inter-network Defense (RID), Trusted Automated eXchange of Indicator Information (TAXII), Simple Object Access Protocol (SOAP), and other standard like reputation services (Repute, DKIM), which fit this category. In the following, we elaborate on the different category of standards and how they are used to automate information sharing within organizations.

## 3.2 A Privacy Risk in Standards

In this section, we highlight the various risks associated with information sharing. For that, *anonymized examples depicted from sharing operations utilizing the standard schemas for information sharing*. For illustration, we label the leaking fields with different colors depending on class of data being exposed, specifically, we designate red color for PII fields, light blue for non-PII sensitive fields (e.g., related to business context), and yellow for inference-leaking fields. This is: for inference, for sensitive, and for PII. In the following, we highlight such risk through various examples obtained from real data.

**IODEF worm report.** An example of a CSIRT reporting an instance of the Code Red worm, encoded in IODEF, is depicted in Figure 1 (notice that a substantial part of the document is omitted, and only essential information is shown for demonstration). As shown, the document contains contact information (name, registry handle, email) for the constituent responsible for the incident report. This type of information may become personally identifiable in the case when contact information of a particular individual is used. The document also includes other fields that are less sensitive. This includes reporting time, record datetime, IP addresses of the node or network that were targeted in the attack, as well as the targeted service port number.

In this example, the Code Red worm attempted to target the HTTP port for a host with an IP address of *192.0.2.200*. The raw HTTP request sent by the worm is captured in the report. The worm intended to fiddle with the web server and the request was presumably an attempt for a buffer overflow attack in order to escalate to administrative privileges. Consequently, if the worm was successful in gaining access to the machine the information captured from the raw HTTP request may become highly sensitive. However, we know from the "assessment" field in the document that it was a failed attempt.

```
1  <IncidentID name="csirt.example.com">189493</
       IncidentID>
2  <ReportTime>2001-09-13T23:19:24+00:00</ReportTime>
3  <Assessment><Completion="failed" type="admin"
       /></Assessment>
4  <ContactName>Example.com CSIRT</ContactName>
5  <Email>contact@csirt.example.com</Email>
6  <Address>192.0.2.200</Address>
7  <Address>192.0.2.16/28</Address>
8  <Service \hlcyan{ip_protocol="6"}>
9  <Port>80</Port>
10 <Expectation action="block-host" />
11 <DateTime>2001-09-13T18:11:21+02:00</DateTime>
12 <Item>GET /default.ida?XXXXXXXXXXXXXXXXXX'</Item>
13 <DateTime>2001-09-14T08:19:01+00:00</DateTime>
14 <Description>sent to constituency-con-
       tact@192.0.2.200</Description>
```

**Figure 1: Annotated private, sensitive, and confidential information (inference) information in IODEF.**

**MAEC package dynamic triage.** An example to demonstrates how a package using the MAEC standard can be used to capture multiple dynamic analysis tool outputs for a malware instance is shown in Figure 2—shortened for summary only. It builds upon static triage example that shows the actions and details of the process tree associated with the instance. As depicted, the packaged output has few fields that may be considered sensitive, such as the domain name of the command & control server (*reallybadguy.com*; only for illustration). Exposing the domain name server to entities outside of the trusted community may inform the attacker about the detection of their malware instance, and thus link the malware reported by the subject with a campaign. In addition to exposing the domain name field, the output also includes a field about bundle actions that are associated with the malware and the status of these actions. In this example, the malware successfully created a file on the host filesystem but failed to resolve the DNS for the command control server. Other countless examples that demonstrate various levels of risk to information that falls under one or more of the categories above exist. By showing those examples above, we hope to trigger interest in the community of pursuing research on understanding the level of leakage (and its context) for various of those sharing schemes in various application contexts.

## 3.3 Privacy Leakage Assessment

With a clear understanding of what constitute attributes that would result in privacy violation, the presence or absence of such attributes in one instance of sharing could be used to assess privacy leakage in total. Informally speaking, we define a **privacy leakage metric**, a single number associated with instances of standards when fully utilized to quantitively analyze the existing (potentially) private information in them. This metric can then be used to obtain numbers for each field in the schema of the standard, that could be aggregated to reflect a single score.

```
1 <FileObj:File_Name>dg003_improve_8080_V132.exe</
2 <cyboxCommon:Type ...>MD5</cyboxCommon:Type>
3 4EC0027BEF4D7E1786A04D021FA8A67F
4 <maecBundle:Action id=maec-example-act-19 context=Host
      action_status=Success>
5 create file</cybox:Name>
6 <maecBundle:Action id=maec-example-act-15 context=Network
      action_status=Fail>
7 <cybox:Name xsi:type="maecVocabs:
      DNSActionNameVocab-1.0">send dns
      query</cybox:Name>
8 <URIObj:Value>4.test.3322.org.cn</URIObj:Value>
```

**Figure 2: MAEC for dynamic triage.**

An example scoring system that assigns real values to various pieces of information in a sharing standard is as follows 1) Score 0: Public data or non-leaking, 2) Score 1: Inferential data, 3) Score 2: Sensitive data, 4) Score 4: PII data. The aggregate score is simply a summation of the scores for each field in the data schema of the standard. The weights on the scoring metric can be modified depending on the level of emphasis on each class of data. In this example, one point is given for an inferential field, two points for a sensitive field, and four points for a PII field. Since PII is protected through regulations, we give more weight for a PII field to be twice as significant as sensitive field and four times as significant as inferential field. Interpreting the score is trivial; the standard is considered highly privacy-leaking when the score number is large and more privacy-friendly when the number is small.

Table 1 shows privacy leaking fields in schemas of various information sharing standards using and following a similar analysis to the one in the previous section. As can be seen, out of the four classes of standards the languages class have leaked the most personally identifiable information fields, particularly CyBox with a score of 56, STIX with a score of 36 and XCCDF with a score of 38. We notice that those standards embody various standards, and are inclusive of a large schema (with a large number of fields and attributes that cover multiple applications). Some standards have leaked fields that are considered sensitive but not PII, such as timestamps, host attributes like IP addresses, and organizational information. Other standards, such as enumeration standards, have mostly inference leaks related to vulnerability indices and platform identification numbers, which could potentially utilized and misused by an adversary.

Notice also that we do not advocate a specific scoring system for the risk assessment, since such scoring is context dependent. For example, an organization that views confidential information, information to do with business-related matters, might score confidential and inferential information higher than PII, since PII is not important to their security operation.

### 3.4 Architectural Solutions for Privacy

A first step towards ensuring privacy is understanding the risk highlighted earlier through the actual sharing paradigms. Using a concrete notion of privacy, it would be then required to provide a technical solutions to meet such privacy notion, while enabling queries on the data shared using the various standards.

Privacy, often addressed through preserving primitives [9], is not the only optimization parameter that could be taken into account, but also the structure of the community of trust. For example, highly homogenous and trusted communities, e.g., a result of public-public partnership, could get away without implementing the partitioned architecture for optimization, but rather using minimization (i.e., for what is being shared, and for how long) on the raw data, thus achieving a higher accuracy, and better efficiency.

Architectural innovation in information sharing is required to improve practicality. Such innovation is facilitated by the differential nature of data and sharing communities, and we argue that they should be taken into to realize efficient sharing solutions. However, to take them into account, further research would be required for understanding the hidden costs in implementing such architecture, the actual trade-off provided by such split architecture, and how to perform complex queries and function (the ultimate purpose of information sharing) on such split architecture, also in a privacy preserving manner.

## 4 QUALITY OF SHARING

So far we have focused on the issue of privacy associated with the sharing, and the threat of sharing due to poorly understood communities of trust, which deserve further considerations. Another important research issue in the context of information sharing for actionable intelligence is the quality of shared information. Without high quality of shared information, no actionable intelligence can be obtained. Unfortunately, this issue is not well understood in the literature, and requires further exploration by identifying meaning of quality, and basic methods and tools for assessing it.

We believe that the quality of indicators is of paramount importance to the end-goal of information sharing: a timely indicator, like a source of attack, could be used to defend against an emerging attack, unlike a stale indicators that could be hardly used for postmortem analysis. Thus quality of indicators is a central issue in information sharing, and requires further attention for realizing the proper definitions, tools for quantification, and incentives for improvement. Little work, however, has been done in the literature on understanding this central notion.

To assess the value and quality of an indicator is a nontrivial task: if a consumer in the information sharing community knew the information provided to him through the sharing community, he would not need the sharing of the data in the first place. One way to deal with the quality of indicators is to use historical information provided by various community members as a metric for their quality [12]. A community member that provided information that turned to be useful and timely in the past could be annotated as a quality indicators provider, and vice versa. However, such approach for determine the quality of indicators would fall short in multiple aspects. First, it assesses providers of indicators, rather than individual indicators. Second, certain community members might be well known for certain indicators, e.g., domain names, and other indicators, e.g., binaries, and taking the average of both indicators contributed by them might penalize them, thus not allowing community members to benefit from the (partially) valuable indicators they provide.

**Table 1: Privacy leaking fields in schemas of various information sharing standards and example risk assessment using the indicated scores for the various leaked attributes. Scores are for illustration only.**

| Standard Category | Standard | Privacy Leak | | | |
|---|---|---|---|---|---|
| | | PII (4) | Sensitive (2) | Inference (1) | Score |
| Enumeration | CVE | | | CVE-ID | 1 |
| | CWE | | | CWE-ID | 1 |
| | CAPEC | | Submission:Source, Organization, Date | Relationship:ViewID, TargetForm, Nature, TargetID | 10 |
| | CCE | | cce:modified_reference | cce:cce_id, cce:platform | 4 |
| | CPE | | cpe:title | cpe:platform_id | 3 |
| Scoring Systems | CVSS | | | | 0 |
| | CWSS | | | | 0 |
| Languages | OVAL | contributor | timestamp, submitted:date, status_change, affected:family, platform, title, description | definition, reference | 20 |
| | XCCDF | Benchmark:metadata, test:identity | cpe2:platform-specification, platform, status, test:organization, test:profile, test:target, test:target-address, test:target-facts, test:target-id-ref, test:start-time, test:end-time, test:fact | affected:family, platform, benchmarkId-Type, resolved, test:authenticated, test:priviledged | 38 |
| | MAEC | CommentType:author | ArtifactObj:Raw_Artifact, maecPackage:Configuration_Parameter, maecPackage:Name, maecPackage:Value, maecBundle:Collections.timestamp, AnalysisType:start_datetime, AnalysisType:complete_datetime, AnalysisType:complete_datetime, AnalysisType:lastupdate_datetime, AnalysisType:Comments, CommentType:timestamp | maecBundle:Action, maecBundle:CVE | 26 |
| | CEE | | time, host, dst, ipv4, ipv6, src, port | status | 15 |
| | IODEF | Contact, IncidentSource | DetectTime, StartTime, EndTime, ReportTime | Assessment, IncidentID, AlternativeID | 19 |
| | STIX | stixCommon:Identity, stixCiqIdentity:Specification, xnl:PersonName, stixCommon:Name, xpil:Address, xpil:ElectronicAddressIdentifier, xpil:ContactNumber | timestamp, xpil:OrganizationInfo, xnl:OrganisationName, xpil:Nationalities/xpil:Country/xal:NameElement | | 36 |
| | Cybox | EmailMessageObj:Recipient, EmailMessageObj:From, AddrObj:Address_Value, EmailMessageObj:Raw_Header, Contributors, ContributorType: Role, Name, Email, Phone | HTTPSessionObj:Value, URIObj:Value, PortObj:Port_Value, ArtifactObj:Raw_Artifact, EmailMessageObj:Date, X509CertificateObj:Subject, X509CertificateObj:Issuer, TimeType: Start_Time, End_Time, Produced_Time, Received_Time, Observation_Location, Observable_Location, ContributorType:Organization, Date, Contribution_Location | | 65 |

## 5 CONCLUDING REMARKS

This paper provides a roadmap of issues that need to be explored in order to realize efficient and effective information sharing paradigms for actionable intelligence. With the evolution of the threat and security landscape, no single defender will be able to defend against all threats alone, calling for the utilization of sharing paradigms. However, in order to utilize such paradigms a finer understanding of the various issues associated with sharing is required, including, but not limited to, the underlying community of trust, threat and use models, and privacy highlighted through measurable contexts from various sharing standards and datasets. We argue that utilizing the differential nature of data and communities of trust could be nicely utilized as a feature for optimizing the overhead of sharing, the role that machine learning could play in understanding and assessing the quality of indicators.

## Acknowledgement

This work was done partly while A. Mohaisen was visiting the Air Force Research Laboratory as a summer faculty fellow, and sponsored by AFOSR. This work is also supported in part by NSF grant CNS-1643207 and the National Research Foundation (NRF) grant NRF-2016K1A1A2912757.

## REFERENCES

[1] —. 2016. Cyber Information Sharing and Collaboration Program. http://www.dhs.gov/topic/cybersecurity-information-sharing. (May 2016).
[2] 2014. *Framework for improving critical infrastructure cybersecurity.* Technical Report. National Institute of Standards and Technology.
[3] R Barnes, B Schneier, C Jennings, T Hardie, B Trammell, C Huitema, and D Borkmann. 2015. *Confidentiality in the Face of Pervasive Surveillance: A Threat Model and Problem Statement.* Technical Report.
[4] Facebook Inc. 2016. ThreatExchange. https://developers.facebook.com/products/threat-exchange/. (May 2016).
[5] Shafi Goldwasser. 1997. Multi party computations: past and present. In *Proceedings of the sixteenth annual ACM symposium on Principles of distributed computing.* ACM, 1–6.
[6] Chris Johnson, Lee Badger, David Waltermire, Julie Snyder, and Clem Skorupka. 2016. *Guide to Cyber Threat Information Sharing.* Technical Report. NIST.
[7] Charles A. Kamhoua, Andrew P. Martin, Deepak K. Tosh, Kevin A. Kwiat, Chad Heitzenrater, and Shamik Sengupta. 2015. Cyber-Threats Information Sharing in Cloud Computing: A Game Theoretic Approach. In *IEEE CSCloud.* 382–389.
[8] Charles A. Kamhoua, Anbang Ruan, Andrew P. Martin, and Kevin A. Kwiat. 2015. On the Feasibility of an Open-Implementation Cloud Infrastructure: A Game Theoretic Analysis. In *8th IEEE/ACM UCC.* 217–226.
[9] Myungsun Kim, Aziz Mohaisen, Jung Hee Cheon, and Yongdae Kim. 2016. Private Over-Threshold Aggregation Protocols over Distributed Datasets. *IEEE Trans. Knowl. Data Eng.* 28, 9 (2016), 2467–2479.
[10] Javvad Malik. 2016. Threat Intelligence Sharing: The Only Way to Combat Our Growing Skills Gap. Information Security Magazine. (May 2016).
[11] Aziz Mohaisen and Omar Alrawi. 2013. Unveiling Zeus: automated classification of malware samples. In *Proc. of ACM WWW.*
[12] Aziz Mohaisen and Omar Alrawi. 2014. AV-Meter: An Evaluation of Antivirus Scans and Labels. In *Proc. of DIMVA.*
[13] Helen Nissenbaum. 2009. *Privacy in context: Technology, policy, and the integrity of social life.* Stanford University Press.
[14] Teri Robinson. 2014. Breaches, malware to cost $491 billion in 2014, study says. http://bit.ly/1gNXu90. (2014).
[15] Julie Ryan. 2012. Use of Information Sharing Between Government and Industry as a Weapon. *Leading Issues in Information Warfare & Security Research* 1 (2012).
[16] Deepak K. Tosh, Shamik Sengupta, Charles A. Kamhoua, Kevin A. Kwiat, and Andrew P. Martin. 2015. An evolutionary game-theoretic framework for cyber-threat information sharing. In *IEEE ICC, London, United Kingdom, June 8-12, 2015.* 7341–7346.
[17] Deepak K. Tosh, Shamik Sengupta, Sankar Mukhopadhyay, Charles A. Kamhoua, and Kevin A. Kwiat. 2015. Game Theoretic Modeling to Enforce Security Information Sharing among Firms. In *IEEE CSCloud.* 7–12.
[18] An Wang, Aziz Mohaisen, Wentao Chang, and Songqing Chen. 2015. Delving into Internet DDoS Attacks by Botnets: Characterization and Analysis. In *Proc. of IEEE DSN.*
[19] An Wang, Aziz Mohaisen, Wentao Chang, and Songqing Chen. 2015. Revealing DDoS Attack Dynamics behind the Scenes. In *Proc. of DIMVA.*
[20] An Wang, Aziz Mohaisen, and Songqing Chen. 2017. An Adversary-Centric Behavior Modeling of DDoS Attacks. In *37th IEEE International Conference on Distributed Computing Systems, ICDCS 2017, Atlanta, GA, USA, June 5-8, 2017.* 1126–1136.