The Trusted Automated eXchange of Indicator Information (TAXII™)

Julie Connolly, Mark Davidson, Matt Richard, Clem Skorupka

11/08/2012

Trademark Information

TAXII and STIX are trademarks of The MITRE Corporation.

This technical data was produced for the U. S. Government under Contract No. HSHQDC-11-J-00221, and is subject to the Rights in Technical Data-Noncommercial Items clause at DFARS 252.227-7013 (NOV 1995)

Feedback

Community input is necessary for the success of TAXII. Feedback on this or any of the other TAXII Specifications is welcome and can be sent to taxii@mitre.org. Comments, questions, suggestions, and concerns are all appreciated.

Executive Summary

Throughout the history of computing and cyber security, the sophistication, speed, and impact of cyber attacks has grown. Today, nuisance-motivated virus outbreaks and internet worms have given way to skillful, persistent adversaries seeking specific information, such as online banking credentials, intellectual property, or national security information. Other nefarious objectives, such as information deception, disruption, and destruction, may also be sought by this new breed of threat actor. Cyber defense strategies must adapt to these new actors and attacks.

Cyber threat information sharing has proven critical in the fight against today's sophisticated cyber adversaries. Today, an ever-growing number of organizations actively share cyber threat data [1] to get a more complete view of adversary activity and help prioritize the organization's own cyber defenses. "My detection becomes your prevention" rings true in these circles. Current cyber threat information sharing, however, is either a time-consuming, manual process or a limited-scope automation effort tied to a particular cyber threat information sharing community or technology. The capability to broadly share a rich set of cyber threat information – beyond such basic elements as IP addresses and file hashes – in an automated manner does not exist today. The Trusted Automated eXchange of Indicator Information (TAXII) effort, a community-driven framework to facilitate cyber threat information sharing, aims to fill this void. Through various technical mechanisms, TAXII seeks to extend indicator sharing to enable robust, secure, and high-volume exchanges of significantly more expressive sets of cyber threat information. This paper describes the TAXII effort, the motivation behind TAXII, its goals, components, and development approach.

TAXII is a set of technical specifications and supporting documentation to enable sharing of actionable cyber threat information across organization and product/service boundaries. TAXII defines protocols and data formats for securely exchanging cyber threat information for the detection, prevention, and mitigation of cyber threats in real time. TAXII is not a specific information sharing initiative or technology, and it does not attempt to define trust agreements, governance, or non-technical aspects of cyber threat information sharing. Instead, TAXII empowers organizations to achieve improved situational awareness about emerging threats, and enables organizations to easily share the information they choose with the partners they choose, while leveraging existing relationships and systems.

Developed with community consensus and participation, TAXII enables efficient and comprehensive cyber threat information exchange through *automation* and *articulation* of a detailed cyber threat information model. To achieve this, TAXII utilizes *a standardized cyber threat information representation* and defines *a supporting exchange framework*. The TAXII model will enable the automatic delivery and receipt of a rich set of cyber threat information to support a diverse set of information sharing needs.

TAXII covers a broad range of use cases, technologies, specifications, and implementations. TAXII is being developed in a phased approach that addresses these use cases in a sequential manner. This strategy aims to deliver an initial set of cyber threat information sharing capabilities, followed by continuing development to increase the breadth of capabilities. TAXII will leverage existing protocols

and specifications wherever possible and integrate with current information sharing mechanisms where appropriate to both reduce implementation costs and allow for rapid adoption by mature organizations that are already sharing information.

TAXII is being developed within an open community forum where participation from individuals and organizations is encouraged and welcomed. MITRE serves as the moderator of the TAXII community on behalf of the Department of Homeland Security (DHS). Mutual exchanges of ideas, collaboration, and the pursuit of technical excellence are key factors for its success. To contribute to the development of TAXII, implement TAXII in a product, use a product that implements TAXII, or just maintain awareness about TAXII, please join the community by visiting http://taxii.mitre.org/ or emailing taxii@mitre.org.

Overview

Cyber threat information sharing has proven critical in the fight against today's sophisticated cyber adversaries. Today's cyber threat information sharing state of the practice, however, is either a time-consuming, manual process or a limited-scope automation effort tied to a particular cyber threat information sharing community or technology. The capability to broadly share a rich set of cyber threat information – beyond such basic elements as IP addresses and file hashes – in an automated manner does not exist today. The community-driven Trusted Automated eXchange of Indicator Information (TAXII), however, aims to fill this void by providing technical mechanisms for cyber threat information sharing that are applicable to a wide range of sharing needs yet flexible enough to accommodate existing cyber threat information sharing implementations. This paper describes the TAXII effort for potential cyber defender adoptees, the motivation behind TAXII, its goals, components, and development approach.

Background

Throughout the history of computing and cyber security, the sophistication, speed, and impact of cyber attacks has grown. Today, nuisance-motivated virus outbreaks and internet worms have given way to skillful, persistent adversaries seeking specific information, such as social security numbers, intellectual property, national security information, and more. Other nefarious objectives, such as information deception, disruption, and destruction, may also be sought by this new breed of threat actor. These adversaries will lie in wait for their lucky break – e.g., someone clicks – until they can obtain the access and information they need. Cyber defense strategies must adapt to these new actors and attacks.

Cyber attacks from such advanced actors are growing in scope and increasing in frequency [2]. Today's predominant patch and block, alert-driven defensive strategy, while effective against some types of threats, fails to stop advanced attacks and provides no knowledge of what an adversary does once the network is penetrated. A more effective framework for thinking about cyber defense is the *cyber kill-chain*, originally conceptualized by Lockheed Martin [3], which is presented in Figure 1.

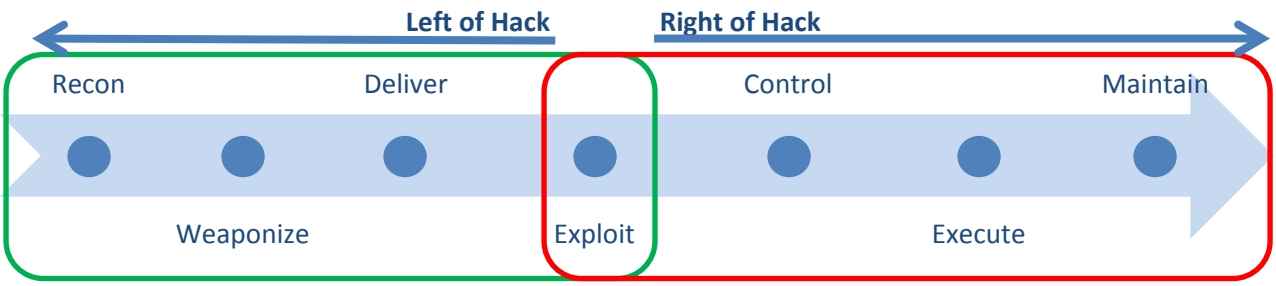

Figure 1. Cyber Kill-Chain

The *cyber kill-chain* breaks down the stages of an advanced cyber attack, in order to augment attack and attacker comprehension and defense. Table 1 describes each cyber kill-chain phase in more detail, to include examples.

Cyber Kill Chain Phase	Description	Example
Recon(naissance)	The adversary identifies and investigates targets.	Web mining against corporate websites and online conference attendee lists
Weaponize	The set of attack tools are packaged for delivery and execution on the victim's computer/network.	The adversary creates a trojanized PDF file containing his attack tools.
Deliver	The packaged attack tool or tools are delivered to the target(s).	The adversary sends a spearphishing email containing the trojanized PDF file to his target list.
Exploit	The initial attack on the target is executed.	The targeted user opens the malicious PDF-file and the malware is executed.
Control	The adversary begins to direct the victim system(s) to take actions.	The adversary installs additional tools on the victim system(s).
Execute	The adversary begins fulfilling his mission requirements.	The adversary begins to obtain desired data, often using the victim system as a launch point to gain additional internal system and network access.
Maintain	Long-term access is achieved.	The adversary has established hidden backdoors on the target network to permit regular re-entry.

Table 1. Cyber Kill Chain Phases

Early steps of the kill-chain, commonly called *left of exploit* or *left of hack*, represent an opportunity to proactively detect and mitigate threats before the adversary establishes a foothold. In the later steps of the kill-chain, commonly called *right of exploit* or *right of hack*, incident detection/response can be exercised along with mission assurance of critical assets. By understanding an adversary's kill-chain, defenders have more opportunity to discover and respond to an attack. Cyber defense strategies need to move to the left of hack in order to anticipate and proactively mitigate threats before they are much harder to find and eradicate using traditional, right of hack detection and response techniques.

Active defense informed by the adversary's kill-chain requires detailed cyber intelligence. Cyber intelligence—or the collecting, analyzing and countering of cyber security threat information—starts with gathering information about attacks, such as spear-phishing email header and content, urls to malicious links, and malware analysis-derived artifacts like Command and Control (C2) domain names and IP addresses. With a corpus of threat data, skilled cyber analysts can group patterns of similar activity, attribute activity to certain threat actors, quickly identify and implement mitigation strategies, and anticipate the launch of similar attacks in the future.

To fully realize the benefits of cyber intelligence, organizations need to share cyber threat data, if not defensive strategies and more, with trusted partners. Analysis across this broader scope paints an even

more compelling picture of adversary activity and necessary defensive actions. As noted in the Security for Business Innovation Council report [1], "Sharing cyber-risk intelligence and defensive strategies has become imperative in today's threat landscape. No organization can realistically sit in isolation and still be able to defend itself." By understanding adversaries' behavior against a range of targets over a period of time, defenders can identify a set of indicators and a robust set of adversary tactics, techniques and procedures (TTPs). By sharing threat information, defenders gain valuable insights into an attacker's overall goals and strategies. This, in turn, improves the defenders' ability to predict attacker behavior and create more dynamic defenses.

Cyber Threat Information Sharing State of the Practice

Today, an ever-growing number of organizations are actively share cyber threat data [1]. Sharing has proven invaluable in painting a richer picture of adversary activity and prioritizing an organization's own cyber defenses. "My detection becomes your prevention" rings true in these circles.

Cyber threat information sharing today takes many forms. It occurs within different sharing communities that use different sharing models, use a range of sharing methods, and exchange a variety of threat data. This section describes each of these aspects in more detail.

Cyber Threat Information Sharing Communities

As the value of cyber threat information sharing has grown, the number and kinds of cyber threat information sharing communities has also grown. Today, there are three primary types of sharing communities:

- Peer
- Commercial, and
- Government

Across all cyber threat information sharing communities, trust is very important. Sharing sensitive cyber threat data could expose the sharing organization to reputation damage, litigation or worse. Sloppy sharing, for instance, could potentially tip off the adversary and render the resulting analytical products useless. Various protections, such as data handling restrictions, data anonymization, NDAs, and the establishment of mutual trust relationships, must be in place [4]. The data handling restrictions become particularly important when organizations engage in more than one cyber threat information sharing community, a common practice used to assemble a more complete picture of relevant adversary activity. In this scenario, what is shared in one community is not necessarily shareable with another community.

Peer communities are the most common, where organizations and/or individuals with a common purpose unite to improve the collective defense against a common adversary or set of adversaries. The industry Information Sharing and Analysis Centers (ISACs), such as the Financial Services ISAC (FS-ISAC) [5] and the Research and Education Network ISAC (REN-ISAC) [6], and regional collaboration groups, such as the Advanced Cyber Security Center (ACSC)[7] in Massachusetts, are examples of formal peer cyber threat information sharing communities. Formal cyber threat information sharing communities

generally have some sort of protections in place, such as non-disclosure agreements (NDAs), to protect the sensitive threat data and the organizations that share it. Less formal cyber threat information sharing communities also exist, often a collection of individuals who know and trust each other and personally vouch for new members. In both cases, the degree of trust among participants usually correlates with the amount and quality of threat information shared.

Commercial communities comprise paid members who are largely anonymous to each other; the commercial entity centrally manages and disseminates the cyber threat information. For instance, iDefense, Symantec, McAfee, Mandiant, Arbor Networks, and others offer different flavors of this kind of cyber threat information sharing service. In some cases, special hardware or software is required that automatically captures and provides various threat data back to a central server for ingest and analysis. In other cases, the threat information is a one-way data feed sent to the customer. Subscribing to one of these services can offer a quick way to jumpstart access to a broad set of cyber threat data. The threat information provided by commercial communities can be broader than what is offered in more specialized peer groups and may not always be applicable to an organization's specific defensive needs.

Finally, government cyber threat information sharing communities are established and managed by the government, are voluntary or mandatory, and include both government and private industry participants. Like the commercial community, the participating organizations and the data they provide are largely anonymous to other participants. The government entity controls threat data collection and dissemination. The Department of Homeland Security's Cyber Information Sharing and Collaboration Program (CISCP), a partnership between DHS and Critical Infrastructure and Key Resources (CIKR) organizations and the Department of Defense (DoD)-Defense Industrial Base (DIB) Collaboration Information Sharing Environment (DCISE) [8], a partnership between the DoD and its DIB contractor community, are two examples.

Cyber Threat Information Sharing Models

Each cyber threat information sharing community typically uses one of three sharing models:

- hub and spoke
- peer-to-peer
- source/subscriber

In the hub-and-spoke model, one entity controls receipt and dissemination of the cyber threat data. The hub entity often anonymizes the collected threat information and provides additional analysis to participants. This model is commonly seen in the commercial and government cyber threat information sharing communities.

In peer-to-peer models, participants share and receive threat data directly with and from other participants. There is no one data owner. When the data is shared, it is shared equally with everyone and usually with the source clearly identified.

The last sharing model is the source/subscriber model. This model is used most frequently by commercial cyber threat information providers. The information provider (source) pushes out regular

cyber threat information feeds to all subscribers. If the subscriber also uses a software or hardware threat component, e.g., anti-virus software, then local subscriber threat data may also be collected and passed back to the source. Often, the source/subscriber cyber threat information is encoded in a proprietary manner and may lack critical context information about individual intrusion attempts. However, this model can offer a quick way for organizations with limited analyst resources to jumpstart access to a broad set of cyber threat data.

Cyber Threat Information Sharing Methods

There are several methods in use for sharing cyber threat information. Often, the method of sharing plays a significant role in the types, volume and nature of information shared within a community. Some exchange mediums place limitations on the type of content that is easily shared, while others actively promote certain types of exchanges. For example, an email listserv will tend to discourage participants from sharing malware samples due to the logistics of passing around files that might be detected by anti-virus tools.

Common methods for exchanging information include:

- Email listserv
- Protected portal discussions
- Wiki / Collaborative Editing
- Data repositories
- Data feed / notification
- Chat / real-time communications

Unfortunately, most of these methods do not lend themselves well to the automated ingestion of cyber threat information. Most consumers routinely take raw cyber threat information from these feeds and synthesize it into their own internal database.

Open architecture, standards-based indicator and incident information sharing are in use. The REN-ISAC, for example, employs a standards-based indicator sharing system called the Collective Intelligence Framework as part of its Security Event System [9]. Argonne National Laboratory, part of the U.S. Department of Energy, developed the Cyber Federated Model (CFM) [10] to share block lists with federated communities of partners. Regional and sector-specific consortia have also developed standards-based models for sharing indicators within their community. However, none of these efforts has resulted in a cross-community standard for interoperable indicator sharing, nor have they developed a means to rapidly establish peer-to-peer, versus hub and spoke, channels for trusted exchange.

Cyber Threat Information Shared

Today, commonly shared cyber threat data are *indicators* of potential nefarious system or network activity—or cyber "observables"[11] of interest, such as IP addresses, domain names, file names, or email addresses. A classic example is that of the IP address "watch list." Organizations routinely create and share lists of IP addresses known or suspected of playing some role in malicious activity, including

hosts serving malicious content, sources of phishing emails or servers acting as command and control (C2) nodes.

In some cases, the shared information is focused on a particular threat, such as botnets, as seen in the PRISEM network in Washington state [12] and Abuse Helper [13] used by many Computer Security Incident Response Teams (CSIRTs). In other cases, numerous details surrounding a particular cyber event, including the malware in use and other TTPs, may be shared. Sometimes, only the malware or malware characteristics are shared. The extent of information shared is generally established by the sharing community and the degree of trust.

Limitations

Cyber threat information sharing has quickly matured the cyber defensive capabilities of numerous organizations. However, current sharing approaches do not realize cyber threat information sharing's full potential. In particular, the majority of sharing is manual, as individual communiqués of self-formatted indicators between small groups of analysts via email, blogs or portals. Since the processes are manual, they are time-consuming, repetitious, and untimely, requiring each organization to re-write or translate the threat information–often also manually–into a large variety of formats. Furthermore, the information is too-often shared via insecure transport, such as unencrypted mailing lists. Because of the variety of data formats and sharing protocols in use, as well as the manual processes involved, cyber threat information sharing does not scale beyond a few trusted organizations.

Another factor that makes cyber threat information sharing inefficient, less scalable, and less timely is the use of proprietary and/or technology-specific formats, necessitating the development of numerous conversion scripts and modules to permit cyber threat information sharing outside of pre-defined communities. For those communities with some degree of automation, their sharing models are generally narrow in scope and use either a commercial or proprietary solution–with the aforementioned conversion issues–or one tailored to its particular community.

Finally, the atomic nature of most cyber indicators is a fundamental limitation. For instance, once a particular IP address is identified as suspect, the effort and therefore cost for the adversary to switch to another IP address is virtually zero. As such, indicator sets such as IP watch lists or other easy to change data elements are said to be highly "perishable." Relying solely on atomic indicators without context can provide numerous false positive alerts and waste significant analyst time.

Motivation

A broader cyber threat information sharing solution is needed, one that spans different communities and sharing models, permits different sharing methods, and supports a wide range of threat data. In particular, the broad goals of the ideal cyber threat information sharing solution are to:

- Permit faster, more accurate cyber threat information sharing
- Reduce tedious human analyst activities—e.g., data entry—and free up analyst time to do the more valuable analysis work
- Move more well understood threats from human analysis to machine processing

- Enable the automated sharing of a wide range of threat data beyond simple, atomic indicators to enable active defense
- Protect the sharing of sensitive threat data
- Permit automatic ingestion of shared threat data into local threat knowledgebases but with context and discretion, thus requiring fewer analyst-eyes needed to screen
- Enable cross organization analyst collaboration on the truly challenging issues

Standardized threat data formats and sharing implementations will achieve these goals. As noted in the *Roadmap to Intelligence-driven Information Security* [1],

> "Automated data-exchange systems need to be established to remove the dependency on specific people. In addition, harmonized standards for representing attack information in machine-readable format, delivering it securely, and consuming it in real time would help to enable automation."

Additionally, as noted in *Breaking Down Barriers to Collaboration in the Fight Against Advanced Threats* [14]:

> "[There is a] Lack of interoperable standards to describe advanced threats – The security industry has yet to align behind a set of uniform, machine-readable standards to capture, integrate and communicate threat information. While some leading-edge organizations have standardized data formats for internal systems, extensive manual processing is still the norm when sharing threat indicators, attack forensics or security intelligence with outside parties. Not only must the global security community harmonize data formats for describing basic incidents, it must also establish consistency in expressing the variability and nuances inherent in advanced threats."

TAXII fills this void.

What Is TAXII?

The Trusted Automated eXchange of Indicator Information (TAXII) is a set of technical specifications and supporting documentation for the secure, platform-independent exchange of high fidelity cyber threat information. TAXII specifications are designed to enhance interoperability of different cyber security solutions rather than espouse a specific technology or product, and vendors are encouraged to incorporate support for TAXII specifications within their cyber security products and services. By supporting TAXII specifications, vendors will enhance the value of their solutions by allowing their customers to leverage actionable intelligence from multiple sources.

Developed through community consensus and participation, TAXII will enable more efficient and comprehensive threat exchange through *automation* and the *articulation* of a detailed, cyber threat information model. To achieve this, TAXII utilizes *a standardized cyber threat information representation* and defines *a supporting exchange framework*. The high level vision for TAXII is shown in Figure 2, streamlining cyber threat information exchange by replacing multiple

proprietary indicator and exchange formats with one set of consensus-defined solutions. By focusing on information representation and secure data exchange, TAXII maximizes the opportunities for solution vendors to integrate TAXII into their cyber security product and service offerings. TAXII offers an agreed-upon way of describing and exchanging machine-consumable cyber threat indicators, leaving vendors free to determine how their products produce, consume, or otherwise take advantage of TAXII-specified data flows.

This section describes the TAXII concept in more detail.

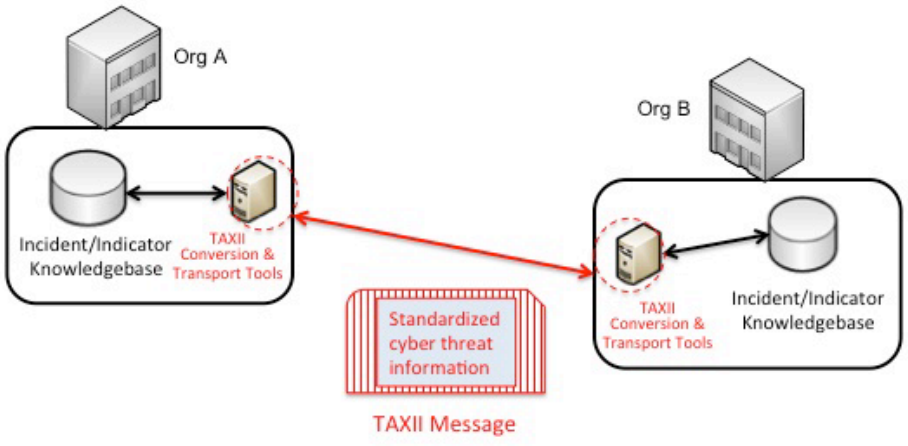

Figure 2. High Level TAXII Vision

Goals of TAXII

The goals of TAXII are to

- Enable *timely and secure* sharing of threat information both within and between cyber defender communities
- Leverage consensus standards to enable the sharing of actionable indicators and more across organization and product/service boundaries
- Extend indicator sharing to enable robust, secure, high-volume exchanges of significantly more expressive sets of cyber threat information
- Support a broad range of use cases and practices common to cyber threat information sharing communities
- Leverage existing mature standards, where appropriate
- Eventual adoption by one or more international standards organizations

TAXII is not creating a sharing community. Rather, it enables communities to share.

TAXII addresses current cyber threat information sharing shortcomings by providing common, open specifications for transporting cyber threat information messages, with capabilities such as encryption, authentication, addressing, alerting, and querying between systems.

Standardized Information Representation

The first part in automating cyber threat information exchange is establishing consensus on **what** is being shared. TAXII uses a standardized language for expressing cyber threat information, called the Structured Threat Information eXpression (STIX™) [15]. STIX is a community-developed language for the specification, capture, characterization and communication of standardized cyber threat information. STIX provides a unified architecture consisting of constructs to support several types of cyber threat information, including:

- Cyber Observables: standardized descriptions of a cyber artifacts or events, represented in the Cyber Observable eXchange (CybOX) language [11]
- Indicators: cyber observables with context
- Incidents: cyber events of interest
- Adversary Tactics, Techniques, and Procedures (TTPs): the tools, attack methods, and attack implementations used by the adversary, to include malware, exploits, tools, infrastructure, targeting, etc.
- Exploit targets: things that get exploited, such as vulnerabilities and weaknesses
- Courses of action: operational responses to a cyber incident (e.g., proactive mitigations, incident response, or vulnerability/weakness remedies)
- Cyber Attack Campaigns: sets of related adversary activity, to include TTPs, indicators, exploit targets, and incidents
- Threat actors: characterizations of adversaries

The current high level STIX representation is shown in Figure 3. To maximize compatibility and ease of adoption, STIX leverages several existing standards, such as CybOX, the Malware Attribute Enumeration and Characterization (MAEC) [16], the Common Vulnerabilities and Exposures (CVE) [17], and the Common Platform Enumeration (CPE) [18]. Included in STIX are information producer tags as well as a flexible data marking structure to provide helpful context and facilitate information handling restrictions. The producer field, included in each STIX construct, identifies the source of the threat information, e.g., analyst name and organization. While the producer field can be used to locally assign confidence to shared data, it may also facilitate implementation of organization-based handling restrictions. The handling field provided within several key STIX constructs, as well as the STIX data marking construct, also permit application of existing data marking schemes, such as the Traffic Light Protocol (TLP) [19], or the development and use of a custom data marking and handling solution.

Breaking STIX down into individual components enhances flexibility. One or more STIX components can be specified and shared, supporting the variability found with unfolding cyber events and different trust relationships. The larger set of constructs also affords a greater amount of context specification, beyond what is available in atomic indicators. The number and aggregation of different types of cyber threat information overcome the limitations of “perishable” indicators by providing a more complete, context-rich set of shareable threat data, sometimes referred to as behavioral indicators.

Initially, TAXII will focus on indicator exchange but plans to eventually support the exchange of additional cyber threat information based upon the constructs available in STIX.

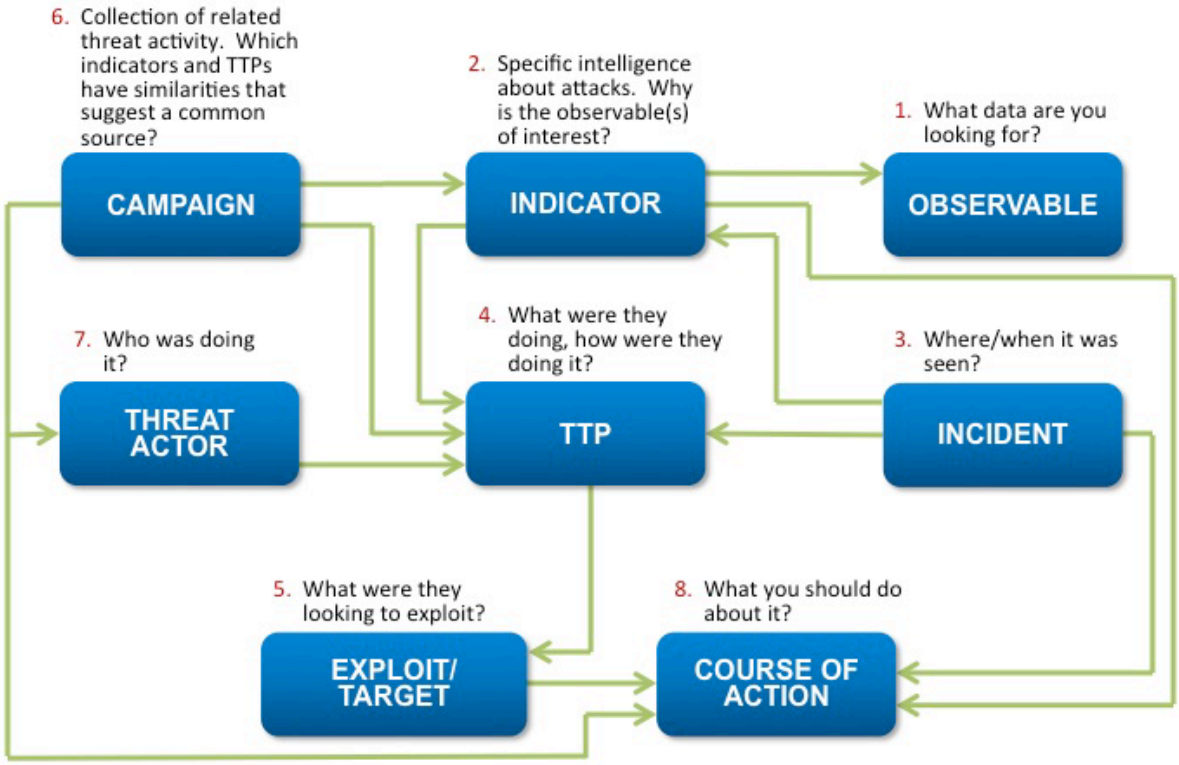

Figure 3. High Level Structured Threat Information eXpression (STIX) representation

An Exchange Framework

The second part in automating cyber threat information exchange is specifying **how** the information is shared. To achieve this, TAXII defines technical specifications and supporting documentation. In particular, TAXII specifications define the set of capabilities needed for successful transport of TAXII messages, or how TAXII messages get from point A to point B. TAXII messages carry a payload of cyber threat data converted into STIX format. The full set of TAXII messages include payload and control messages.

TAXII leverages existing protocols and specifications wherever possible and integrate with current information sharing mechanisms to both reduce implementation costs and allow for rapid adoption by mature organizations that are already sharing information. TAXII is being developed in a modular fashion to support a variety of exchange mechanisms and data formats. Core TAXII concepts are defined independent of the implementation details. Community consensus will drive development priorities of TAXII agents to support various protocol bindings, with the development of more commonly used protocols such as HTTP and SMTP occurring first.

The TAXII exchange framework also ensures that TAXII messages are exchanged securely. Encryption or other security implementation mechanisms draw from the underlying protocol, as appropriate, and will be detailed in the protocol-specific binding specification.

TAXII Use Cases

TAXII is being developed to support common use cases for cyber threat information sharing. The following enumerates these use cases and the high-level requirements that TAXII implements.

Public Alerts or Warnings

These are warnings to the general public, such as various CSIRT advisories, that are broadcast or published to all subscribers. Examples of these include:

- CSIRT advisories
- New malware or exploits in the wild

These alerts are of such a broad nature that no encryption or special access or authorization is required. A digital signature however is important to ensure authenticity. Entity / organization addressing is required to identify the source of the Alert, as well as the recipients for mailing-list-style delivery. The addressing scheme should be scalable, so organizational / entity name-space will need to be defined.

Private Alerts and Reports

Private alerts are similar to Public alerts, except the threat information shared is sensitive and restricted to sharing partners. Sharing partners may be part of a defined, formal cyber threat information sharing community, such as an ISAC, or informal trust relationships with a peer organization. Data delivery mechanisms or services for private alerts should be similar to the public alerts. Since the alerts and reports are presumed sensitive and not for general consumption, it is important that TAXII support appropriate forms of encryption, authentication, authorization, and data tagging. Depending on the nature of the communiques, explicit handling guidelines or markings for restrictions on data sharing may be necessary. A standard set of machine-readable representations for common marking scenarios, to include “Not for Public Release”, “Case Pending: Immediate Responders Use only”, and “Do not share outside of Community”, may be used.

Alerts are generally short, standardized messages with very specific indicators or actions specified, such as “Block this URL that is distributing Malware” or “Apply Patch #12345 from Vendor X”.

Reports are lengthier messages, and they can include incident reports, malware analysis, threat analysis, or other observations that are of a less immediate context.

Query support

It is not uncommon for Threat Analysts to seek information from others within or outside of their sharing community.

RFI: One form of message may be a simple Request for Information (RFI) such as “do you have any information on this piece of malware, this IP address”. It is expected that these types of requests would be handled manually, perhaps landing in a human analyst’s work queue.

Repository Search: For this type of query, it is expected that given organizations will offer up a searchable repository, which may be TAXII / STIX compatible.

Bulk transfer

Various cyber threat information sharing organizations will at times have new members, possibly member companies or other organizations themselves. The new member may require a "data dump" of the organization's threat data repository. Therefore it is expected that a "bulk transfer" mode may need to be supported by TAXII.

These use cases will be supported in an iterative manner, as TAXII development unfolds. Additional information about TAXII development can be found in the section that follows.

TAXII Development Approach

TAXII covers a broad range of use cases, technologies, specifications, and implementations. The TAXII Strategy is a phased plan that will address use cases in a sequential manner.

The TAXII strategy has three phases: Initial Planning and Design, Initial Use Case Build Out, and Sequential Build Out. This strategy aims to deliver capabilities across TAXII in a focused manner, followed by continuing development to increase the breadth of capabilities.

The first phase, Initial Planning and Design, calls for the overall design of TAXII. The work from this phase will provide a framework for the entirety of the TAXII work.

The Initial Use Case Build Out phase calls for the development of each TAXII component, but only as required to support an initial use case. This phase is characterized by focused development of aspects of TAXII that support an initial use case. During this phase, the work done in the Initial Planning and Design phase will be evaluated.

The last phase, Sequential Build Out, calls for incremental development of TAXII. During this phase TAXII will mature as support for additional use cases is added.

Components of TAXII

As shown in Figure 4, TAXII is a set of technical specifications and supporting documentation. Each TAXII component specification defines a portion, or supported use case, of TAXII. The suite of technical specifications that comprise TAXII are enumerated and described in this section.

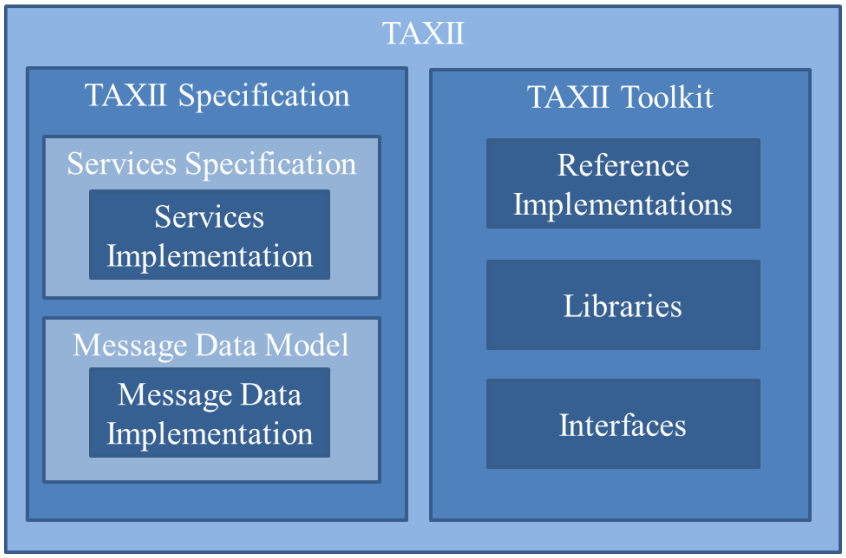

Figure 4. TAXII Representation

TAXII Specification

The TAXII Specification defines component specifications and provides guidance and requirements for how those specifications interoperate within TAXII. The TAXII Specification will also include a set of supported use cases.

TAXII Service Specification

The TAXII Service Specification defines the services that must be implemented in order to be considered TAXII compliant. The Service Specification describes information exchange at a high level; it does not bind services to any particular exchange mechanism (e.g., HTTP, SMTP, SOAP).

TAXII Services Implementations

TAXII Services Implementations bind the TAXII Service Specification to particular exchange mechanisms (e.g., HTTP, SMTP, SOAP). Each Services Implementation provides the technical guidance and requirements for implementing the Services Specification in a particular exchange mechanism.

TAXII Message Data Model

The TAXII Message Data Model defines the structure of TAXII messages, including the header, payload, control, and data messages. The Message Data Model utilizes STIX™ for the payload of TAXII messages. The Message Data Model does not directly bind the data model to any particular format.

TAXII Message Data Implementations

A TAXII Message Data Implementation binds the TAXII Message Data Model, including the STIX payload, to a particular format (e.g., XML, protobuf, etc.). Each Message Implementation defines the technical guidance and requirements for using a particular format to express the Message Data Model.

The TAXII Toolkit

The TAXII Toolkit is provided to support adoption of TAXII and assist in the development of compatible capabilities. The toolkit includes a collection of reference implementations, a set of tools, and a collection of libraries and interfaces, which will evolve over time based upon the contributions and needs of the community. This section describes the planned areas of development in support of TAXII.

Reference Implementations

For particular components identified in the TAXII Components Specification, a reference implementation may be provided. A reference implementation is software developed to demonstrate the TAXII specification and enable testing. A reference implementation may be accompanied by a reference architecture document to provide guidance for other possible implementations of TAXII.

Utilities

A number of utilities will be developed to support TAXII. Such utilities may include:

- utilities to consume an artifacts (i.e., email messages) and create TAXII formatted data,
- utilities to exercise components of a transport specification binding, and
- utilities to represent TAXII data in a human-readable fashion.

Libraries and Interfaces

Libraries and interfaces will be developed to support actions that are frequently performed within an application context. Such libraries and interfaces may include a library to create and manipulate TAXII formatted data or a library to exchange TAXII messages by a specific transport specification binding.

Summary

We have presented TAXII as a community-driven framework to address the emerging needs for cyber threat information sharing, including automation, security, consistency, richness of expression, and interoperability. TAXII is intended to be a set of guidelines that will encourage increased sharing of threat information, and as such incorporates many existing protocols and mechanisms in use by the different communities. Initial TAXII development will focus on essential features such as encryption and alert support, with the intent of eventually supporting a comprehensive set of cyber threat information sharing use cases.

DHS, MITRE and the rest of the TAXII community welcome your participation in defining and implementing TAXII. If you would like to contribute to the development of TAXII, implement TAXII in a product, use a product that implements TAXII, or just maintain awareness about TAXII, please join the community.

Website: taxii.mitre.org
Email: taxii@mitre.org

References

[1] "Getting Ahead of Advanced Threats: *Achieving Intelligence-Driven Information Security,* Recommendations from Global 1000 Executives," Security for Business Innovation Council, January 2012.

[2] "Advanced Persistent Threats: A Decade in Review," Command Five Pty Ltd, June 2011. URL http://www.commandfive.com/papers/C5_APT_ADecadeInReview.pdf

[3] E.M. Hutchins, M.J. Cloppert and R.M Amin PH.D., "Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains," *Proc. 6th Int'l Conf. Information Warfare and Security* (ICIW 11), Academic Conferences Ltd., 2010, pp. 113–125; URL http://www.lockheedmartin.com/content/dam/lockheed/data/corporate/documents/LM-White-Paper-Intel-Driven-Defense.pdf

[4] D.F. Vasques, O.P. Acosta, S. Brown, E. Reid, and C. Spirito, "Conceptual Framework for Cyber Defense Information Sharing within Trust Relationships," *Proc. 4th Int'l Conference on Cyber Conflict* (CYCON 2012), NATO CCD COE Publications, 2012, pp. 1-17; URL http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6243990

[5] The Financial Services Industry Sharing and Analysis Center (FS-ISAC). URL http://www.fsisac.com/

[6] The Research and Education Networking Information Sharing and Analysis Center. URL http://www.ren-isac.net/

[7] The Advanced Cyber Security Center (ACSC). URL http://www.massinsight.com/initiatives/cyber_security_center/

[8] The Department of Defense (DoD)-Defense Industrial Base (DIB) Collaboration Information Sharing Environment (DCISE). URL http://www.dc3.mil/dcise/

[9] The Security Event System's Collective Intelligence Framework (CIF), The Research and Education Network Information Sharing and Analysis Center (REN-ISAC). URL http://code.google.com/p/collective-intelligence-framework/

[10] Argonne National Laboratory Cyber Fed Model (CFM). URL http://web.anl.gov/it/cfm/

[11] Cyber Observable eXpression (CybOX). URL http://cybox.mitre.org/

[12] The Public Regional Information Security Event Management (PRISEM) Effort. URL http://www.cyber.st.dhs.gov/experimentsandpilots/

[13] AbuseHelper. URL http://www.enisa.europa.eu/activities/cert/support/chiht/tools/AbuseHelper/

[14] "Breaking Down Barriers to Collaboration in the Fight Against Advanced Threats," RSA, February 2012.

[15] Barnum, Sean. "Standardizing Cyber Threat Intelligence Information with the Structured Threat Information eXpression (STIX™)," MITRE Corporation, July 2012.

[16] The Malware Attribute Enumeration and Characterization (MAEC). URL http://maec.mitre.org/

[17] Common Vulnerabilities and Exposures (CVE). URL http://cve.mitre.org/

[18] Common Platform Enumeration (CPE). URL http://cpe.mitre.org/

[19] Traffic Light Protocol (TLP). URL http://www.us-cert.gov/tlp/