THE MITRE CORPORATION

Sample Usage of TAXII

Version 1.0 (draft)

Mark Davidson, Charles Schmidt

11/16/2012

The Trusted Automated eXchange of Indicator Information (TAXII™) specifies mechanisms for exchanging structured cyber threat information between parties over the network. This document walks through detailed examples of how TAXII might be used by members of a sharing community.

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.

Table of Contents

1 Introduction

Trusted Automated eXchange of Indicator Information (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 requirements are provided through multiple related specifications (see Section 1.1). This document provides a detailed descriptions of how sharing communities could use TAXII for their interactions.

The purpose of this document is to illustrate the use of TAXII 1.0 in common usage scenarios using concrete examples. As such, this document does not contain normative requirements for the implementation of TAXII.

Readers of this document should read it in conjunction with the TAXII Services Specification [1], the TAXII HTTP Binding Specification [2], and the TAXII Message Binding Specification for XML [3] as this document makes references to all three documents. Readers are advised to read these specification first, and then use this document to clarify expected behavior within each of the TAXII specifications.

1.1 TAXII Specifications

TAXII is defined by multiple, interrelated specifications. This section describes the specifications that define TAXII.

Services Specification - The TAXII Services Specification provides requirements that govern TAXII services and exchanges. It does not provide details on data formatting or how TAXII messages are transported over a network - such details and requirements can be found in the Protocol Binding Specifications and Message Binding Specifications.

Protocol Binding Specification - Protocol Binding Specifications define the requirements for transporting TAXII messages over the network. There may be multiple Protocol Binding Specifications created for TAXII. Each Protocol Binding Specification defines requirements for transporting TAXII messages using some network protocol (e.g., HTTP). They provide requirements about how the TAXII Services are supported by these network protocols.

Message Binding Specification - Message Binding Specifications define the requirements for representing TAXII messages in a particular format. There may be multiple Message Binding Specifications created for TAXII. Each Messaging Binding Specification defines a binding for TAXII messages (e.g., XML). They provide detailed guidance about how the information in the TAXII messages, as defined in the Services Specification, is actually expressed.

Figure 1 shows how these specifications relate to each other.

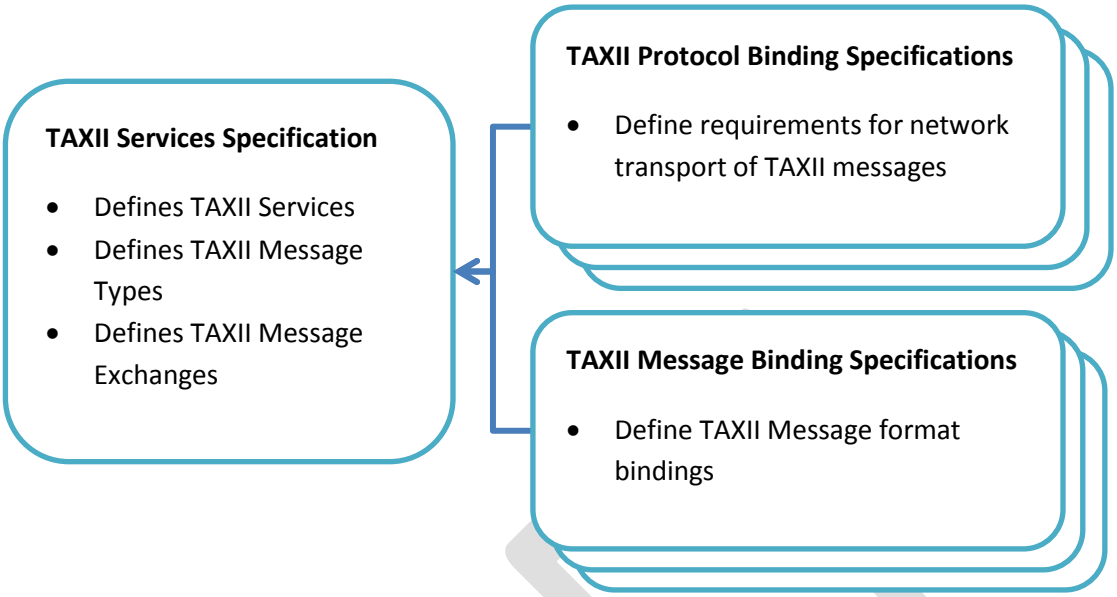

Figure 1 - TAXII Specification Hierarchy

Separation of the Services Specification, Message Binding Specifications, and Protocol Binding Specifications exists to support flexibility as TAXII evolves. Threat information sharing communities often have specific constraints on the types of protocols they are able to support. Rather than binding TAXII to a specific protocol that excludes portions of the community, TAXII's core concepts (i.e., its services and exchanges) are defined separately from the protocol-level support for those concepts. When there is evidence of significant community interest in new protocol and message bindings, TAXII can define support for those bindings without changing its core components.

Two groups that use the same network protocol and message bindings will be capable of automated exchanges of structured threat information. The sharing policies of the participants can limit these exchanges as needed, but the use of compatible TAXII services ensures that whatever sharing is permissible by policy can be effected by the TAXII mechanisms. Groups that use different protocol or message bindings for TAXII will not be able to communicate directly with each other, but because they are still using TAXII Messages and Services at the core of their communications means that it is possible to create gateways that will allow interaction to occur.

1.1.1 STIX

TAXII is designed to support the sharing of structured cyber threat information. The structuring of this information is provided by the Structured Threat Information eXpression (STIX™). STIX is "a collaborative community-driven effort to define and develop a standardized language to represent structured cyber threat information." [4]

This specification does not provide details about the underlying structures defined in the STIX specification, apart from noting that all cyber threat information transported by TAXII is expressed in

"STIX documents". Those interested in learning more about STIX are directed to the STIX web site at https://stix.mitre.org/.

1.2 Terms and Definition

This section defines terms that are assigned a specific meaning within all TAXII specifications:

1.2.1 TAXII Concepts

These terms are used throughout the document to define concepts central to definition of TAXII.

TAXII Data Feed - A collection of structured cyber threat information expressible in one or more STIX documents that can be exchanged using TAXII. All TAXII Data Feeds MUST be assigned a name that uniquely identifies them on a given Producer. Individual pieces of cyber threat information within a TAXII Data Feed are labeled with a timestamp and may have other labels at the producer's discretion.

TAXII Message - A discrete block of information that is passed from one entity to another. A TAXII Message represents either a request (e.g., “Can I subscribe to this TAXII Data Feed?”) or a response (e.g., “Yes.”).

TAXII Message Exchange - A defined sequence of request and response TAXII Messages undertaken by two parties to accomplish a specific activity.

TAXII Service - Functionality hosted by some entity that is accessed or invoked through the use of one or more TAXII Message Exchanges.

1.2.2 TAXII Roles

TAXII Roles are used to denote participants in TAXII according to their high-level objectives in the use of TAXII Services.

Producer - The role of an entity (e.g., a person, organization, agency, etc.) that is the source of structured cyber threat information.

Consumer - The role of an entity that is the recipient of structured cyber threat information.

1.2.3 TAXII Network Components

These terms are used to define the components of a TAXII Implementation using a typical client-server model. Note that these should be considered orthogonal to the TAXII Roles previously defined: An entity might both host a TAXII Server and use a TAXII Client in their role as a TAXII Consumer. The defined network components represent a network-centric view of TAXII participants while the defined roles represent an activity-centric view.

TAXII Server - A TAXII implementation that provides one or more TAXII services. To support this functionality, it is assumed that a TAXII Server is persistently listening for new TAXII network traffic.

TAXII Client - A TAXII implementation that initiates an exchange with a TAXII Server. A TAXII Client does not need a persistent connection on the network to operate but can open a connection when it wishes to interact with a TAXII server and disconnect from the network when this interaction has concluded.

2 An Example of TAXII Support for Source/Subscriber Sharing

Previous work [5] has identified three models used by cyber threat information sharing communities to exchange threat information. These models are:

- Source/subscriber - The information provider pushes out regular information to all subscribers
- Peer-to-peer - Participants share and receive threat data directly
- Hub and spoke - One entity controls receipt and dissemination of cyber threat data that might be collected from multiple sources

This section considers the first of these sharing models as it is the simplest and because this model can be a component of the other two. In a Source/Subscriber model information flow is unidirectional, from a single Source to each Subscriber. Figure 3 shows an example of a Source/Subscriber relationship.

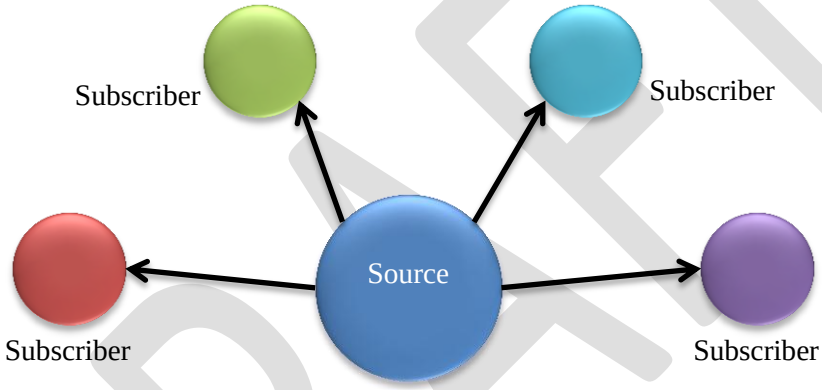

Figure 2 - The Source/Subscriber Sharing Model

For this example, it is assumed all participants support the TAXII HTTP Protocol Binding 1.0 [2] and TAXII XML Message Binding 1.0 [3].

2.1 The Participants

In this model there is a Source that provides cyber threat information to some number of Subscribers who each enter into agreements with this Source to receive periodic content updates. The Subscribers are likely to be unknown to each other, so this represents a set of bi-lateral agreements that each Subscriber makes with the single Source. In this model, the Source is a TAXII Producer, while the Subscribers are TAXII Consumers.

For the purposes of this example, this document assumes the following participants.

Source - A vendor that collects, edits, and then sells threat intelligence information. This vendor might use TAXII to receive threat intelligence information, but for the sake of this example, only its interaction with customers is considered. Customers purchase contracts to receive threat intelligence information.

Contracts are tiered - those in the highest tiers receive the greatest amount of information, while those who purchase lower-tiered contracts get progressively less detail. The vendor also offers a free feed with very limited information as a public service/means to drive business.

Subscriber - A customer of the aforementioned vendor. This customer represents a company that seeks to improve its security posture by making use of available cyber threat information.

2.2 Initial Preparation

Prior to a TAXII exchange, both the Source and all would-be Subscribers must implement TAXII. Given that these participants have different activities in this exchange, their TAXII implementations will have different requirements. The two participants will also need to have an understanding of cyber threat information that is compatible with that information's expression using STIX.

2.2.1 Architectures

Both participants must implement TAXII in order to participate in a TAXII exchange. The key distinction between these implementations is in the TAXII Services they support.

2.2.1.1 Source Architecture

The Source might make use of several TAXII Services (as defined in the TAXII Service Specification). Note that it is possible for the Source to have a TAXII implementation that includes only some of the described TAXII Services. The applicable TAXII Services appear below along with their use and how the Source might provide the equivalent functionality without (or in addition to) using the identified TAXII Services.

2.2.1.1.1 TAXII Discovery Service

Implementing a TAXII Discovery Service would allow a Subscriber to use TAXII to learn which services the Source offers and how to contact them. This might be especially useful if the Source implemented multiple TAXII Feed Management Services with different responsibilities and/or if the Source implemented one or more TAXII Poll Services. Interaction with the Discovery Service would inform the Subscriber as to which TAXII Services were offered, which protocol and message bindings those services supported, how to contact those services, and additional information, such as whether one needed a certain contract tier to access the service in question.

If the Source does not provide a Discovery Service, this information needs to be provided through other means. For example, the Source could host a web page that contained a table with the relevant information. Alternatively, the Source might provide its customers with an application to interact with its services that was pre-configured with the relevant information. Other means of conveying the relevant information to the Subscriber through non-automated means are also possible.

2.2.1.1.2 TAXII Feed Management Service

A TAXII Feed Management Service allows Subscribers to discover what TAXII Data Feeds are offered by the Source, including descriptive information about the nature of those feeds and what forms of delivery the Source is capable of providing for their content. In addition, through the Feed Management

Service a Subscriber could use TAXII Messages to create and manage subscriptions to those TAXII Data Feeds.

If the Source does not provide a Feed Management Service, then advertisement of TAXII Data Feeds and management of subscriptions need to be handled through other means. A list of TAXII Data Feeds might be provided on a web page along with their delivery mechanisms. Subscription to a feed might be automatically integrated to the vendor's purchasing software, automatically creating a subscription when a customer purchases the appropriate contract. Other out-of-band mechanisms for creating subscriptions could also be implemented.

2.2.1.1.3 TAXII Poll Service

A TAXII Poll Service allows Subscribers to pull TAXII Data Feed content instead of having the Source push this content to them. This can be useful to a Subscriber that, for policy reasons, prohibits the establishment of inbound connections. This also allows the Subscriber to collect information at times that are convenient to them rather than waiting for distributions by the Source.

If the Source does not provide a Poll Service then all TAXII Data Feed content must be pushed to an Inbox Service hosted by the Subscriber. Alternatively, content could be disseminated through mechanisms not associated with TAXII.

2.2.1.1.4 TAXII Client for a Subscriber's Inbox Service

If the Source pushes TAXII Data Feed content to the Subscriber, the Source will need a TAXII Client that sends this content to the Subscriber's TAXII Inbox Service. This push functionality does not require a TAXII Service, as defined in the TAXII Services Specification, but does require TAXII-compatible functionality on the part of the Source . Specifically, this functionality would need to format TAXII Data Feed content as a TAXII Message and then make a connection to the appropriate Subscriber Inbox Service to deliver this content.

If the Source does not implement such a TAXII Client, the Source would need to host a Poll Service and all Subscribers would need pull content from the Source. Alternatively, content could be disseminated through mechanisms not associated with TAXII.

2.2.1.2 Subscriber Architecture

Depending on how the Subscriber will be receiving content, they may or may not need to implement a TAXII Service. Specifically, if they are to receive content via push messaging, they will need to host a TAXII Inbox Service, but if they will only be pulling content from the Source the Subscriber does not require a TAXII Service although it would need a TAXII Client. The Subscriber will also need TAXII Clients capable of interacting with some or all of the other TAXII Services the Source provides.

2.2.1.2.1 TAXII Inbox Service

If the Subscriber is to receive TAXII Messages via push messaging, then they must implement a TAXII Inbox Service. The Inbox Service can receive pushed TAXII Data Feed content for the Subscriber. The Inbox Service can also receive subscription alerts informing the Subscriber of current or impending changes to their existing subscriptions.

If the Subscriber does not implement a TAXII Inbox Service then they will need to contact a Poll Service on the Source in order to collect TAXII Data Feed content via pull messaging. Alternatively, they could receive content using mechanisms not associated with TAXII.

2.2.1.2.2 TAXII Client(s)

The Subscriber would need a TAXII Client for any interactions they wished to have with any of the TAXII Services hosted by the Source, such as the Discovery Service, Feed Management Service, or Poll Service. The actual implementation of this functionality could take the form of a single application capable of interacting with all of these services, multiple applications each capable of interacting with a single TAXII Service, or something in between. The Subscriber's TAXII Client(s) would interact with these TAXII Services to learn the TAXII Services offered by the Source, create or manage subscriptions to TAXII Data Feeds, and/or collect TAXII Data Feed content via pull messaging, respectively.

If the Source provides means to accomplish the activities associated with a service via mechanisms other than TAXII, a Subscriber could use those mechanisms instead of implementing the appropriate TAXII Client.

2.2.2 Data Handling - TAXII Data Feeds

In addition to the described architectural components, both participants in a TAXII exchange must share a particular understanding of the structured cyber threat information that TAXII exchanges. As noted earlier, both participants must be able to express and understand structured cyber threat information using the STIX language. However, in addition to this, there must be an understanding of a TAXII Data Feed.

2.2.2.1 Source

TAXII content dissemination is tightly bound to the concept of a TAXII Data Feed. When a subscription is created, it is to a specific, named TAXII Data Feed. As such, all content that the Source makes available via TAXII subscriptions must be assigned to one or more TAXII Data Feeds.

The assignment of content to one or more TAXII Data Feeds is at the Source's discretion. For example, the Source might create a "Full" feed which included all content, and then several additional feeds which included (possibly overlapping) subsets of all content based on criteria such as contract level, membership in other sharing groups, or other special relationships. It is not necessary for TAXII Data Feeds to be the sole means by which content dissemination is controlled. For example, instead of having separate TAXII Data Feeds for each contract level, a vendor might define separate TAXII Data Feeds by the type of threat they discuss and dynamically elide information in those feeds just prior to sending them to a customer based on that customer's contract level. In fact, one could imagine a Source that had only a single TAXII Data Field that all of its customers used where that Source performed a series of rule-based edits to each piece of content prior to delivery so that each customer only saw content that was appropriate to their contract and access rights. That said, most Sources will probably find some benefit to defining some grouping of their content into TAXII Data Feeds as a means of managing access and allowing customers to receive material that they deem relevant to their operations.

In addition to mapping content to one or more TAXII Data Feed, each piece of content must be assigned a timestamp within each TAXII Data Feed to which it is mapped. It should be emphasized that this timestamp exists to provide a label used by exchanges associated with the TAXII Poll Service. It does not need to represent an actual chronological time - the Source could simply label the "first" piece of content in a TAXII Data Feed with the time 0000-00-00T00:00:01 and increment this value with each subsequent piece of content in the TAXII Data Feed. This said, most Sources will probably find it easier to assign timestamps that are, in fact, chronologically meaningful. For example, a piece of content's timestamp could represent the time the content was created, the time the content was disseminated, the time of the event associated with the content, or some other meaning. Further note that the same piece of content could be labeled with a different timestamp in different TAXII Data Feeds. There is also no prohibition against multiple pieces of content within a single TAXII Data Feed sharing the same timestamp label. Since TAXII Data Feed timestamps are only utilized by exchanges with the TAXII Poll Service, if the Source does not provide Poll Service access to the TAXII Data Feed, it would not need to apply timestamp labels to that TAXII Data Feed's content.

2.2.2.2 Subscriber

The Subscriber needs to be aware of the Feed Names associated with TAXII Data Feeds as these names are used when establishing subscriptions and polling for content. Likewise, if the Subscriber uses the Source's Poll Service, they will likely wish to track timestamps associated with prior Poll Requests in order to avoid pulling the same content multiple times. Note that a response to a Poll Request includes the bounding timestamps for the range of TAXII Data Feed content it is returning, but otherwise the Subscriber has no means of learning the timestamp a Source assigned to a particular piece of content. The timestamp is not intended to provide the Subscriber with a precise handle to specific pieces of content but simply a way to request non-overlapping ranges of data.

Once a Subscriber has received a piece of content, there is no further need to associate that content with a specific Source-assigned TAXII Data Feed.

2.3 A Sample Source-Subscriber Interaction

This section looks at the specific TAXII Message Exchanges that could be used to support interactions between a Source and a Subscriber. For this example, it is assumed that both participants have TAXII implementations that include all of the TAXII Services and TAXII Client functionality for their respective roles as described in the previous sections. It is also assumed that the Source has arranged their content into a set of TAXII Data Feeds and assigned timestamps to all this content.

Figure 3 shows a sample sharing arrangement with a single Source and multiple Subscribers. The following sections look at each of the TAXII Message Exchanges that occur between the Source and a single Subscriber in the course of establishing a subscription (steps 1-4) and then having the Subscriber receive TAXII Data Feed content associated with that subscription (step 5). Note that step 3 represents activities outside the scope of the TAXII specifications.

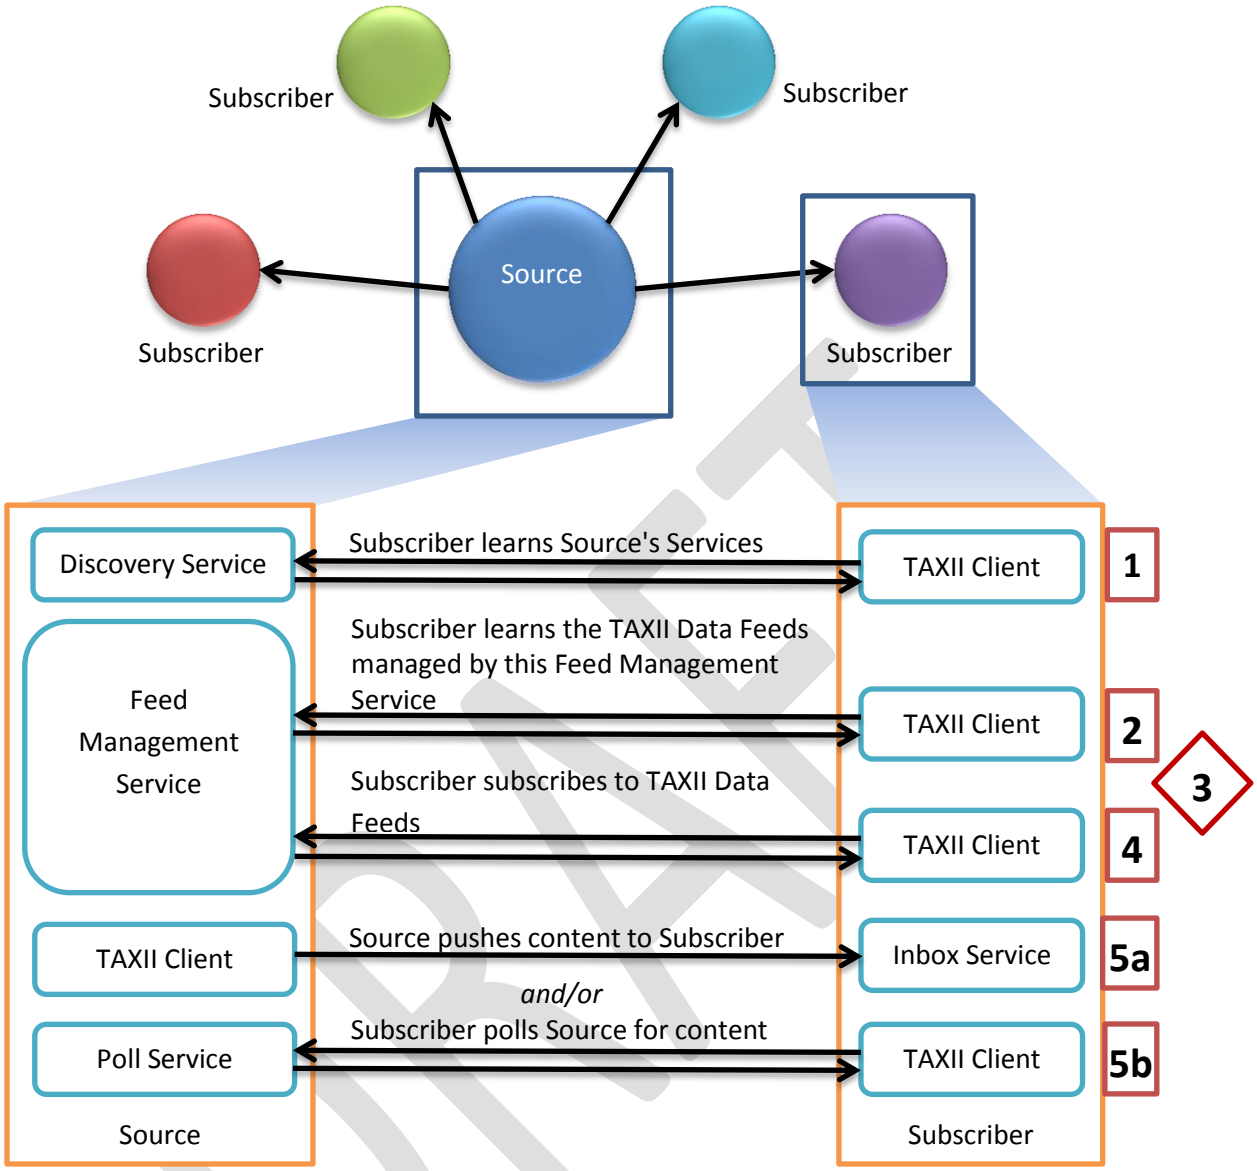

Figure 3 - Source/Subscriber Sharing Model

For the sake of this example, it is assumed the Source and Subscriber have had no preceding interaction and, as such, have effectively no initial knowledge of each other. The one exception is that it is assumed the Subscriber is aware of the existence of the Source's Discovery Service and knows how to contact it. The information necessary to support this interaction might be posted on the Source's web site or in some other public location.

The following sections show what the messages that constitute the identified exchanges might look like using the TAXII HTTP Protocol Binding and the TAXII XML Message Binding. A green background denotes the HTTP request or status lines along with the HTTP headers. A blue back denotes the HTTP message body, if present, expressed using the TAXII XML Message Binding. If a single line is too long to fit on one line of the diagram, a "▶▶▶" is used to indicate line wrapping for readability.

2.3.1 Step 1: Discovery Exchange

As noted above, it is assumed that the Subscriber is aware of the Source's Discovery Service and possesses the information necessary to establish a connection with it. Such necessary information would include the TAXII Protocol Binding to use, the Discovery Service's network address using this binding, and the TAXII Message Binding used by the Discovery Service.

The Subscriber constructs a Discovery Request Message (see Figure 4) and sends it to the Source's Discovery Service. The Source responds with a Discovery Response Message (see Figure 5). For these (and all following) examples, an arbitrary Message ID has been assigned to each message. Note that the In Response To field in the Discovery Response Message is identical to the Message ID field in the Discovery Request Message. Note also that the Source and Subscriber use different formats for their Message ID - this is permitted and should not cause conflicts.

```
GET /?message_type=discovery_request&message_id=▶▶▶
  ▶▶▶a847%2098be%20ffe8%2040ee%203722%20c893%20288b%20d375 ▶▶▶
  ▶▶▶HTTP/1.1

Host: taxiiserver.company.com
Accept: application/xml
Content-Type: application/xml
User-Agent: TAXII client application
X-TAXII-Accept: TAXII_1.0/TAXII_XML_BINDING_1.0
X-TAXII-Content-Type: TAXII_1.0/TAXII_XML_BINDING_1.0
X-TAXII-Protocol: TAXI_HTTP_BINDING_1.0
```

Figure 4 - TAXII Discovery Request Message

```
HTTP/1.1 200 OK

Content-Type: application/xml
Content-Length: 2474
Date: Thu, 15 Nov 2012 08:12:31 GMT
X-TAXII-Content-Type: TAXII_1.0/TAXII_XML_BINDING_1.0
X-TAXII-Protocol: TAXI_HTTP_BINDING_1.0
```

```
<?xml version="1.0" encoding="UTF-8"?>
<TAXII_DiscoveryResponse
    xmlns="http://taxii.mitre.org/messages/xml/1"
    message-id="984987DE-FEAA-95FF-894C-CBCC432988839023"
    in-response-to="a847 98be ffe8 40ee 3722 c893 288b d375">
```

Continued

Continued

```
    <service-instance service-type="DISCOVERY" service-version="TAXII_1.0">
        <protocol-binding>TAXII_HTTP_BINDING_1.0</protocol-binding>
        <message-binding>TAXII_XML_BINDING_1.0</message-binding>
        <service-address>http://taxiiserver.company.com/</service-address>
        <message>Corporate TAXII Discovery Service. This service provides
             information about all our products accessible through
             TAXII.</message>
    </service-instance>
    <service-instance service-type="FEED-MANAGEMENT"
             service-version="TAXII_1.0">
        <protocol-binding>TAXII_HTTP_BINDING_1.0</protocol-binding>
        <message-binding>TAXII_XML_BINDING_1.0</message-binding>
        <service-address>http://taxiiserver.company.com/feeds/public
        </service-address>
        <message>Used to for free public data feeds</message>
    </service-instance>
    <service-instance service-type="POLL" service-version="TAXII_1.0">
        <protocol-binding>TAXII_HTTP_BINDING_1.0</protocol-binding>
        <message-binding>TAXII_XML_BINDING_1.0</message-binding>
        <service-address>http://taxiiserver.company.com/feeds/public
        </service-address>
        <message>To support polling of public data feeds.</message>
    </service-instance>
    <service-instance service-type="FEED-MANAGEMENT"
             service-version="TAXII_1.0">
        <protocol-binding>TAXII_HTTPS_BINDING_1.0</protocol-binding>
        <message-binding>TAXII_XML_BINDING_1.0</message-binding>
        <service-address>https://taxiiserver.company.com/feeds/pay
        </service-address>
        <message>Used to for for-pay data feeds. To purchase a
             subscription, go to http://www.company.com/sales and
             follow the on-screen instructions.</message>
    </service-instance>
    <service-instance service-type="POLL" service-version="TAXII_1.0">
        <protocol-binding>TAXII_HTTPS_BINDING_1.0</protocol-binding>
        <message-binding>TAXII_XML_BINDING_1.0</message-binding>
        <service-address>https://taxiiserver.company.com/feeds/pay
        </service-address>
        <message>To support polling of for-pay data feeds.</message>
    </service-instance>
</TAXII_DiscoveryResponse>
```

Figure 5 - TAXII Discovery Response Message

For this example, it is assumed that the Source hosts a single Discovery Service and two instances each of a Feed Management Service and Poll Service, one pair of which is dedicated to providing for-pay services while the other pair provides free services. The Source might have additional TAXII Services that are only revealed to certain authenticated parties, but since the Subscriber has had no prior interaction with the Source, it cannot be usefully authenticated and those "secret" TAXII Services would not be revealed in this exchange. Note that some TAXII Services share an address (i.e., have the same URL) while others have different addresses. The mapping of Service implementations to network addresses is entirely at the discretion of the Source.

From this exchange, the Subscriber learns of all the Source's TAXII Services (at least all the TAXII Services the Source is willing to tell the Subscriber about), the Protocol and Message Bindings those services support, and how to contact them. The source uses the Message fields in its Discovery Response to describe the different Services.

2.3.2 Step 2: Feed Information Exchange

Having learned of the Source's offered TAXII Services, the Subscriber wishes to learn which Data Feeds are available. In this case, the Subscriber is interested in the for-pay feeds. As such, the Subscriber contacts the Source's for-pay TAXII Feed Management Service with a Feed Information Request Message (see Figure 6) to query it about the available TAXII Data Feeds. The Source responds with a Feed Information Response (see Figure 7). In this response, the Source identifies three TAXII Data Feeds. The Source uses the Description fields associated with these TAXII Data Feeds to describe what these feeds provide. It also notes how content associated with these TAXII Data Feeds can be disseminated. Note that the first two TAXII Data Feeds can be polled for information but the third cannot.

```
GET /feeds/pay?message_type=discovery_request&message_id=▶▶▶
   ▶▶▶a847%2098be%20ffe8%2040ee%203722%20c893%20288b%20d376▶▶▶
   ▶▶▶HTTP/1.1

Host: taxiiserver.company.com
Accept: application/xml
Content-Type: application/xml
User-Agent: TAXII client application
X-TAXII-Accept: TAXII_1.0/TAXII_XML_BINDING_1.0
X-TAXII-Content-Type: TAXII_1.0/TAXII_XML_BINDING_1.0
X-TAXII-Protocol: TAXI_HTTPS_BINDING_1.0
```

Figure 6 - TAXII Feed Information Request Message

```
HTTP/1.1 200 OK

Content-Type: application/xml
Content-Length: 1464
Date: Thu, 15 Nov 2012 09:11:43 GMT
X-TAXII-Content-Type: TAXII_1.0/TAXII_XML_BINDING_1.0
X-TAXII-Protocol: TAXI_HTTP_BINDING_1.0
```

```
<?xml version="1.0" encoding="UTF-8"?>
<TAXII_FeedInformationResponse
    xmlns="http://taxii.mitre.org/messages/xml/1"
    message-id="984987DE-FEAA-95FF-894C-B845503EEF3319AC"
    in-response-to="a847 98be ffe8 40ee 3722 c893 288b de376">
    <feed feed-name="Platinum">
        <description>Our most comprehensive data feed containing up-to-
            the-minute threat information and professional analysis of
            impact and context. </description>
        <delivery-method>TAXII_HTTP_BINDING_1.0</delivery-method>
        <delivery-method>POLL</delivery-method>
        <message-binding>TAXII_XML_BINDING_1.0</message-binding>
        <content-binding>STIX_XML_1.0</content-binding>
    </feed>
    <feed feed-name="Gold">
        <description>Up-to-the-minute threat information covering global
            threats but with no additional analysis.</description>
        <delivery-method>TAXII_HTTP_BINDING_1.0</delivery-method>
        <delivery-method>POLL</delivery-method>
        <message-binding>TAXII_XML_BINDING_1.0</message-binding>
        <content-binding>STIX_XML_1.0</content-binding>
    </feed>
    <feed feed-name="Silver">
        <description>A daily digest of the day's threats.</description>
        <delivery-method>TAXII_HTTP_BINDING_1.0</delivery-method>
        <message-binding>TAXII_XML_BINDING_1.0</message-binding>
        <content-binding>STIX_XML_1.0</content-binding>
    </feed>
</TAXII_FeedInformationResponse>
```

Figure 7 - TAXII Feed Information Response Message

From this exchange, the Subscriber learns of the TAXII Data Feeds that can be managed through this particular Feed Management Service. Since, in our example, the Source hosts two Feed Management Services, the Subscriber might wish to contact the Source's other Feed Management Service next to see what feeds are offered there.

2.3.3 Step 3: Acquire Authorization (Outside TAXII)

The Subscriber identifies the second of the Source's three for-pay TAXII Data Feeds as meeting their operational needs. Access to this TAXII Data Feed's content requires authorization. In this case,

authorization is achieved by paying the Source money using the web interface indicated in the Message field in the Feed Information Response. The Subscriber goes to this web site and provides a credit card number. At the same time, the Source and Subscriber would need to arrange for a way for the Subscriber to authenticate when communicating with the Source. For this example, it is assumed that the Source has the Subscriber establish a username and password combination for this purpose. Alternatively, they could have used certificate-based authentication mechanisms or some other means of authentication provided both parties supported it.

2.3.4 Step 4: Subscription Management Exchange

The Subscriber now seeks to establish a subscription to their chosen TAXII Data Feed. The Subscriber contacts the same Feed Management Service on the Source and sends a Manage Feed Subscription Request message with an action of SUBSCRIBE (see Figure 8). The Subscriber provides the relevant authentication information, transported using the enveloping HTTP Request. The Source validates the authentication information and that the authenticated ability is permitted to subscribe to the given TAXII Data Feed and then responds with a Managed Feed Subscription Response (see Figure 9).

```
POST /feeds/pay HTTP/1.1

Host: taxiiserver.company.com
Accept: application/xml
Content-Type: application/xml
Content-Length: 578
User-Agent: TAXII client application
X-TAXII-Accept: TAXII_1.0/TAXII_XML_BINDING_1.0
X-TAXII-Content-Type: TAXII_1.0/TAXII_XML_BINDING_1.0
X-TAXII-Protocol: TAXI_HTTPS_BINDING_1.0
```

```
<?xml version="1.0" encoding="UTF-8"?>
<TAXII_SubscriptionManagementRequest
    xmlns="http://taxii.mitre.org/messages/xml/1"
    message-id="a847 98be ffe8 40ee 3722 c893 288b d377">
    <feed-name>Gold</feed-name>
    <action>SUBSCRIBE</action>
    <subscription>
        <delivery-method>TAXII_HTTPS_BINDING_1.0</delivery-method>
        <message-binding>TAXII_XML_BINDING_1.0</message-binding>
        <content-binding>STIX_XML_1.0</content-binding>
        <send-to>https://www.customer.com/taxii/inbox</send-to>
    </subscription>
</TAXII_SubscriptionManagementRequest>
```

Figure 8 - TAXII Subscription Management Request Message

```
HTTP/1.1 200 OK

Content-Type: application/xml
Content-Length: 758
Date: Wed, 21 Nov 2012 11:12:31 GMT
X-TAXII-Content-Type: TAXII_1.0/TAXII_XML_BINDING_1.0
X-TAXII-Protocol: TAXI_HTTPS_BINDING_1.0
```

```
<?xml version="1.0" encoding="UTF-8"?>
<TAXII_SubscriptionManagementResponse
    xmlns="http://taxii.mitre.org/messages/xml/1"
    message-id="984987DE-FEAA-95FF-894C-CCB893840004CF67"
    in-response-to="a847 98be ffe8 40ee 3722 c893 288b d377">
    <feed-name>Gold</feed-name>
    <message>Thank you for your business. Your subscription is paid
          up for 6 months.</message>
    <subscription subscription-id="customer.com-JGK8H84BF">

        <delivery-method>TAXII_HTTPS_BINDING_1.0</delivery-method>
        <message-binding>TAXII_XML_BINDING_1.0</message-binding>
        <content-binding>STIX_XML_1.0</content-binding>
        <send-to>https://www.customer.com/taxii/inbox</send-to>
    </subscription>
</TAXII_SubscriptionManagementResponse>
```

Figure 9 - TAXII Subscription Management Response Message

If there was a problem with the Source's request, the Source would respond with a TAXII Error Message. This could include the situation where validating the Subscriber's request would take longer than the timeout of the enveloping HTTP session. In this case, the Error Message type would be PENDING and be accompanied by a timestamp indicating when the Subscriber could repeat their request. At that time, the Source estimates they will have determined if the request is valid and can have a ready response when the Subscriber repeats their request.

Note that the Source's response assigns a Subscription ID to this successfully-requested subscription. This Subscription ID would be used in subsequent requests to manage this subscription, as well as with any attempts to poll to receive subscription content.

For this example, the Subscriber seeks to have the Source deliver TAXII Data Feed content to the Subscriber's Inbox Service using push messaging.

2.3.5 Step 5a: Data Push Exchange

When the Source wishes to provide content associated with a TAXII Data Feed, it uses a TAXII Client to send a STIX Message (see Figure 10) to each Subscriber's Inbox Service. Note that, in our example, the Subscriber authenticates to the Source by providing a username and password. The Inbox Service cannot do this using an HTTP envelope since it is not initiating the connection. As such, the Source would need to trust that the Subscriber correctly provided the address of their Inbox Service. If certificate-based authentication was used, a TLS handshake could authenticate the identity of the Inbox Service.

```
POST taxii/inbox HTTP/1.1

Host: www.customer.com
Accept: application/xml
Content-Type: application/xml
Content-Length: 344
User-Agent: TAXII Feed Pusher 1000 (A TAXII Client)
X-TAXII-Accept: TAXII_1.0/TAXII_XML_BINDING_1.0
X-TAXII-Content-Type: TAXII_1.0/TAXII_XML_BINDING_1.0
X-TAXII-Protocol: TAXI_HTTPS_BINDING_1.0
```

```
<?xml version="1.0" encoding="UTF-8"?>
<TAXII_STIXMessage
    xmlns="http://taxii.mitre.org/messages/xml/1"
    message-id="984987DE-FEAA-95FF-894C-B896830488DE7D6A">
    <subscription-id>customer.com-JGK8H84BF</subscription-id>
    <content-binding>STIX_XML_1.0</content-binding>
    <STIX xmlns="http://stix.mitre.org"/>
</TAXII_STIXMessage>
```

Figure 10 - TAXII STIX Message

Note that the content itself might be changed or elided by the Source to reflect data dissemination policies. For example, certain information might be limited to dissemination to members of particular communities and, if the Subscriber had not demonstrated membership, this information would need to be omitted from any push to the Subscriber. For the sake of this example, the actual STIX content has been removed since the details of the STIX document are outside the scope of TAXII.

The Subscriber's Inbox Service receives the STIX Message and the STIX content is passed along to the relevant mechanisms that process structured cyber threat information.

2.3.6 Step 5b: Feed Poll Exchange

If the Subscriber had chosen to use polling rather than have the Source push content to their Inbox Service, the Subscriber would need to utilize a Feed Poll Exchange to retrieve content. Alternately, some Sources might allow Subscribers to poll on any established Subscription as a means of retrieving older

data or to retrieve new data before a Source's regular pushes sent it out. If the Subscriber wished to poll the Source, it would send a Poll Request Message (see Figure 11) to the Source's Poll Server. Two timestamp fields bound the range of material to be collected in response to a Poll Request Message. A Subscriber might choose, for their first Poll Request, to omit these fields, thus requesting all material currently in the Source's TAXII Data Feed. For this example, the beginning timestamp is provided, indicating a lower bound to the collection, but the second timestamp is omitted, indicating a lack of an upper bound. While a Source controls the exact meaning of the timestamps used in its TAXII Data Feeds, a conventional interpretation of this would be, "provide all content newer than the given time." The Subscriber would need to provide authentication information with this request.

```
POST /feeds/pay HTTP/1.1

Host: taxiiserver.company.com
Accept: application/xml
Content-Type: application/xml
Content-Length: 336
User-Agent: TAXII client application
X-TAXII-Accept: TAXII_1.0/TAXII_XML_BINDING_1.0
X-TAXII-Content-Type: TAXII_1.0/TAXII_XML_BINDING_1.0
X-TAXII-Protocol: TAXI_HTTPS_BINDING_1.0
```

```
<?xml version="1.0" encoding="UTF-8"?>
<TAXII_PollRequest xmlns="http://taxii.mitre.org/messages/xml/1"
    message-id="a847 98be ffe8 40ee 3722 c893 288b d377">
    <feed-name>Gold</feed-name>
    <begin-timestamp>2012-11-07T00:00:00Z</begin-timestamp>
    <subscription-id>customer.com-JGK8H84BF</subscription-id>
</TAXII_PollRequest>
```

Figure 11 - TAXII Poll Request

```
HTTP/1.1 200 OK

Content-Type: application/xml
Content-Length: 518
Date: Mon, 10 Dec 2012 11:52:22 GMT
X-TAXII-Content-Type: TAXII_1.0/TAXII_XML_BINDING_1.0
X-TAXII-Protocol: TAXI_HTTPS_BINDING_1.0
```

```
<?xml version="1.0" encoding="UTF-8"?>
<TAXII_PollResponse xmlns="http://taxii.mitre.org/messages/xml/1"
    message-id="984987DE-FEAA-95FF-894C-800F7EACA647A6A"
    in-response-to="a847 98be ffe8 40ee 3722 c893 288b d377">
    <begin-timestamp>2012-11-07T00:00:00Z</begin-timestamp>
    <end-timestamp>2012-11-12T14:43:34Z</end-timestamp>
    <subscription-id>customer.com-JGK8H84BF</subscription-id>
    <content-binding>STIX_XML_1.0</content-binding>
    <STIX xmlns="http://stix.mitre.org"/>
</TAXII_PollResponse>
```

Figure 12 - TAXII Poll Response

The Source's Poll Service receives this request, verifies the authentication information and that the poll request is valid. The Source then collects the indicated information and formats it into a Poll Response Message (see Figure 12), which is then returned to the Subscriber. Note that the Poll Response Message includes a Begin Timestamp value that matches the corresponding field in the Poll Request Message (indicating that the Source used the requested lower bound for its data collection) but also includes an End Timestamp value. The latter indicates the upper bound of the returned collection. When the Subscriber makes subsequent poll requests, they can use this timestamp, incremented by a minimal value (e.g., `2012-11-12T14:43:34.001Z`), as the lower bound to ensure that they do not pull the same content twice but that they also do not miss any provided content.

3 An Example of TAXII Support for Peer-to-Peer Sharing

This section provides a very simple example of how a minimal implementation of TAXII might be used by a community of users utilizing a Peer-to-Peer sharing model. In a Peer-to-Peer model, individual community members share directly with each other rather than having all content disseminated from a central clearing house. Figure 13 shows a sample Peer-to-Peer sharing model

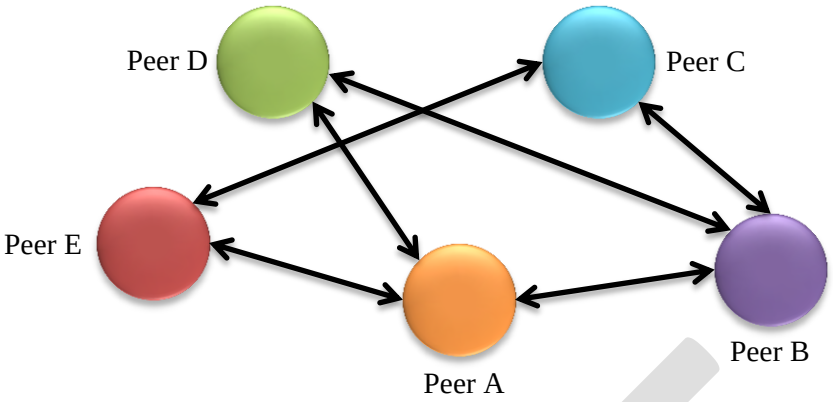

Figure 13 - Peer-to-Peer Sharing Model

In this example, it is assumed that a sharing community has an agreement as to what information each member will share. To minimize infrastructure requirements for each Peer in the group, the decision is made to host a single Discovery Service for the entire sharing community. Every Peer in the community, using procedures outside the scope of TAXII, registers an Inbox Service with this Discovery Service. When any community member has new information to share with the community, that member uses the Discovery Service to get the list of all Inbox Services within the community and then sends a STIX Message containing the new information to each of these Inbox Services. This use of TAXII obviates any need for TAXII Data Feeds or Feed Management Services - instead, community members understand that any new information automatically gets shared to every Inbox Service registered on the communal Discovery Service.

3.1 The Participants

In this model all Peers serve both as TAXII Producers and Consumers, sending and receiving content. In this particular example, there is also a Discovery Service hosted somewhere. All community members would know of this Discovery Service and have access to it. (For privacy reasons, the Discovery Service would likely require all Discovery Requests to be authenticated.) This Discovery Service could be hosted by one of the Peers within the community or it could be an independent entity. In this example, the latter is assumed.

3.2 Initial Preparation

Prior to a TAXII exchange, all Peers in the sharing community must implement TAXII. The community, however, made the decision to minimize the impact of this requirement. As such, all Peers must implement an Inbox Service as well as a TAXII Client capable of interacting with the community's central Discovery Service as well as all community Inbox Services. There are no other requirements on a TAXII implementation that this sharing arrangement imposes. Specifically, there is no need for any of the individual participants to implement Feed Management Services or Poll Services, and only the single instance of the Discovery Service is required.

Moreover, this sharing arrangement does not require any understanding of TAXII Data Feeds. Since all community members automatically share with all other community members, there is no need to subscribe to TAXII Data Feeds or even to organize data according to TAXII Data Feeds.

3.3 A Sample Peer-to-Peer Interaction

This section looks at the specific TAXII Message Exchanges that could be used to support interactions between Peers in this sharing community. For this example, it is assumed that the identified Peers are already community members and that all community members have an entry on the Discovery Service that identifies their Inbox Service. How prospective members might join such a sharing community is beyond the scope of TAXII.

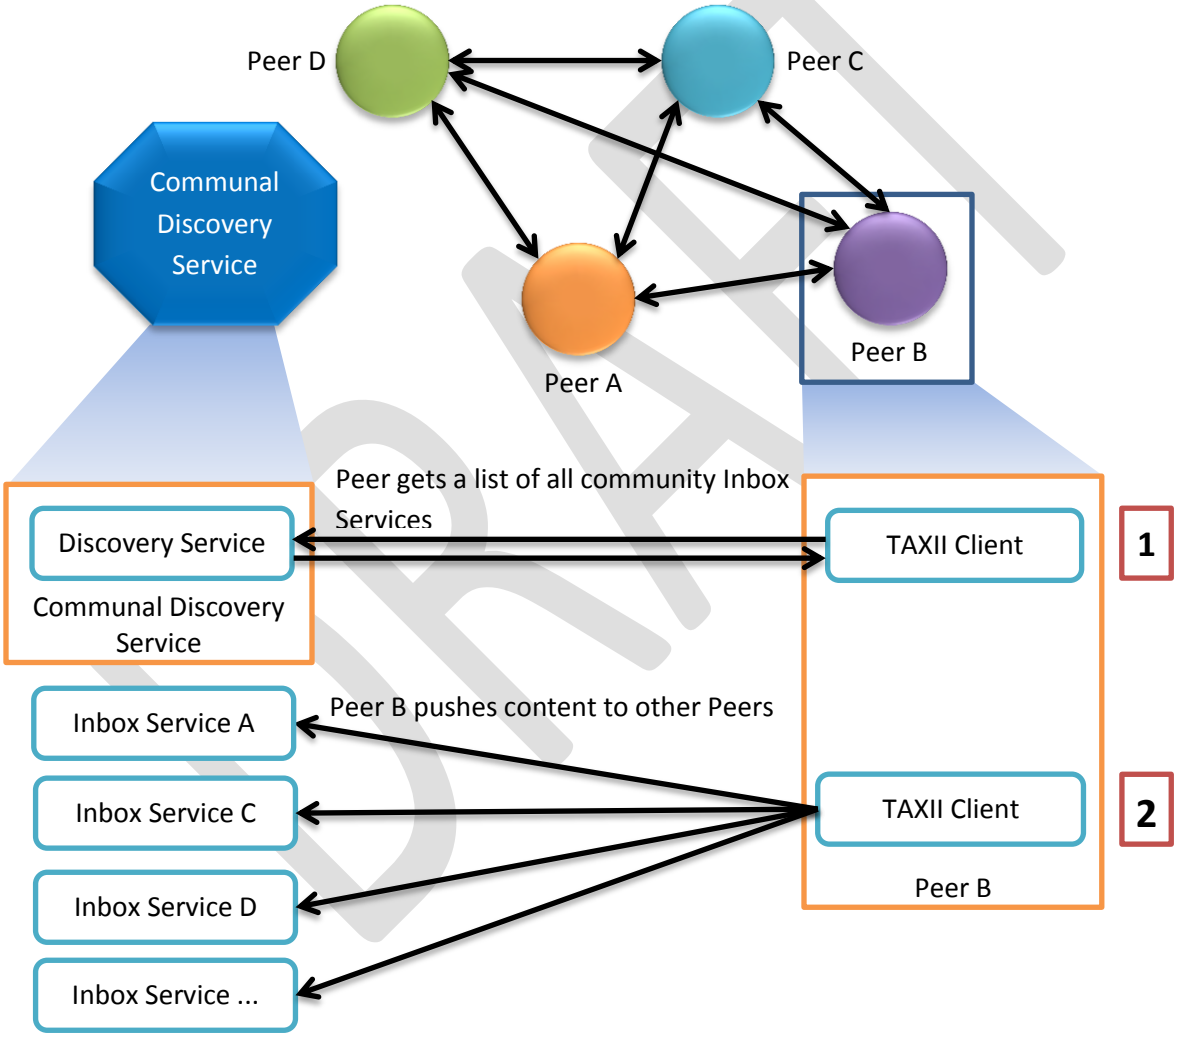

Figure 14 - Example Peer-to-Peer Sharing Exchange

Figure 14 shows an example of how information would be shared using this model. In this example, Peer B develops some cyber threat information it wishes to share with the other members of its sharing community. It first contacts the communal Discovery Service and from it receives a list of current Inbox

Services for the sharing community. It then initiates a Data Push Exchange with each of these Inbox Services, pushing the new content to each Peer in the community.

The following sections look at the messages of both of these sets of exchanges.

3.3.1 Step 1: Discovery Exchange

In this exchange, Peer B contacts the communal Discovery Service to collect latest set of contact information for the Inbox Services for the members of the sharing community. (See Figure 15) The communal Discovery Service responds with a list of the Inbox Services of community members. (See Figure 16)

```
GET /TaxiiDiscovery/?message_type=discovery_request▶▶▶
  ▶▶▶&message_id=a84798beffe840ee3722c893288bd375 HTTP/1.1

Host: serviceregistry.sharinggroupcoordinator.com
Accept: application/xml
Content-Type: application/xml
User-Agent: Sharegroup Client App
X-TAXII-Accept: TAXII_1.0/TAXII_XML_BINDING_1.0
X-TAXII-Content-Type: TAXII_1.0/TAXII_XML_BINDING_1.0
X-TAXII-Protocol: TAXI_HTTPS_BINDING_1.0
```

Figure 15 - TAXII Discovery Request Message

```
HTTP/1.1 200 OK

Content-Type: application/xml
Content-Length: 2865
Date: Wed, 14 Nov 2012 08:22:13 GMT
X-TAXII-Content-Type: TAXII_1.0/TAXII_XML_BINDING_1.0
X-TAXII-Protocol: TAXI_HTTPS_BINDING_1.0
```

```
<?xml version="1.0" encoding="UTF-8"?>
<TAXII_DiscoveryResponse
    xmlns="http://taxii.mitre.org/messages/xml/1"
    message-id="984987DE-FEAA-95FF-894C-CBCC432988839023"
    in-response-to="a84798beffe840ee3722c893288bd375">
    <service-instance service-type="INBOX" service-version="TAXII_1.0">
        <protocol-binding>TAXII_HTTP_BINDING_1.0</protocol-binding>
        <message-binding>TAXII_XML_BINDING_1.0</message-binding>
        <content-binding>STIX_XML_1.0</content-binding>
        <service-address>http://taxii.peerA.com/</service-address>
        <message>ABC Corp.</message>
    </service-instance>
    <service-instance service-type="INBOX" service-version="TAXII_1.0">
        <protocol-binding>TAXII_HTTP_BINDING_1.0</protocol-binding>
        <message-binding>TAXII_XML_BINDING_1.0</message-binding>
        <content-binding>STIX_XML_1.0</content-binding>
        <service-address>http://www.peerB.com/taxii/inbox</service-address>
        <message>BobCom</message>
    </service-instance>
    <service-instance service-type="INBOX" service-version="TAXII_1.0">
        <protocol-binding>TAXII_HTTP_BINDING_1.0</protocol-binding>
        <message-binding>TAXII_XML_BINDING_1.0</message-binding>
        <content-binding>STIX_XML_1.0</content-binding>
        <service-address>http://www.peerC.com/svcs/taxii</service-address>
        <message>Corp.com</message>
    </service-instance>
    <service-instance service-type="INBOX" service-version="TAXII_1.0">
        <protocol-binding>TAXII_HTTPS_BINDING_1.0</protocol-binding>
        <message-binding>TAXII_XML_BINDING_1.0</message-binding>
        <content-binding>STIX_XML_1.0</content-binding>
        <service-address>https://peerD.com/inbox</service-address>
        <message>DigitalData Corp.</message>
    </service-instance>
    ...
</TAXII_DiscoveryResponse>
```

Figure 16 - TAXII Discovery Response Message

Note that, unlike the previous example, the Discovery Response Message does not list the Discovery Service itself. What to include in a Discovery Response Message is at the discretion of the information

Provider. In this case, the sharing community felt that such information was unnecessary to include in the response.

3.3.2 Step 2: Data Push Exchanges

At the conclusion of the Discovery Exchange, Peer B has the latest contact information for all of its Peers in the sharing community. It now contacts each of these Inbox Services and pushes the new content to each. Figure 17 shows an example of one such message. Note that in this message Peer B adds additional TAXII Message header fields (Not HTTP headers - though that is allowed as well) - a sharing community might agree on a set of additional TAXII header fields for their own internal data management.

```
POST / HTTP/1.1

Host: taxii.peerA.com
Accept: application/xml
Content-Type: application/xml
Content-Length: 529
Date: Wed, 14 Nov 2012 08:25:32 GMT
X-TAXII-Accept: TAXII_1.0/TAXII_XML_BINDING_1.0
X-TAXII-Content-Type: TAXII_1.0/TAXII_XML_BINDING_1.0
X-TAXII-Protocol: TAXI_HTTP_BINDING_1.0
```

```
<?xml version="1.0" encoding="UTF-8"?>
<TAXII_STIXMessage
    xmlns="http://taxii.mitre.org/messages/xml/1"
    xmlns:share="http://sharinggroupcoordinator.com/ext-hdr"
    message-id="a84798beffe840ee3722c893288bd378">
    <extended-headers>
        <share:source>Peer B</share:source>
        <share:label>Green</share:label>
    </extended-headers>
    <subscription-id>customer.com-JGK8H84BF</subscription-id>
    <content-binding>STIX_XML_1.0</content-binding>
    <STIX xmlns="http://stix.mitre.org"/>
</TAXII_STIXMessage>
```

Figure 17 - TAXII STIX Message

4 Conclusion

This document has provided one possible example of how TAXII can be used to support models for the sharing of structured cyber threat information. Actual sharing scenarios can vary significantly, but the services offered by TAXII should be able to provide useful building blocks from which virtually any sharing model can be supported. Agreement as to the behavior of these building blocks should allow a

greater degree of information sharing to be accomplished, and the use of standardized structured formats and exchange protocols allows for a greater degree of automation than is possible in most sharing communities today.

5 Bibliography

[1] M. Davidson and C. Schmidt, "The TAXII Services Specification Version 1.0," The MITRE Corporation, Bedford, MA, 2012.

[2] M. Davidson and C. Schmidt, "The TAXII HTTP Binding Specification Version 1.0," The MITRE Corporation, Bedford, MA, 2012.

[3] M. Davidson and C. Schmidt, "The TAXII Message Binding Specification for XML Version 1.0," The MITRE Corporation, Bedford, MA, 2012.

[4] The MITRE Corp., "STIX - Structured Threat Information Expression," 1 October 2012. [Online]. Available: https://stix.mitre.org/. [Accessed 19 October 2012].

[5] U.S. Department of Homeland Security, "Trusted Automated eXchange of Indicator Information (TAXII ™)," U.S. Department of Homeland Security, Washington D.C., 2012.

[6] Defense Advanced Research Projects Agency, "RFC 793 - Transmission Control Protocl," The Internet Engineering Task Force, 1981.

[7] J. Postel, "RFC 768 - User Datagram Protocol," The Internet Engineering Task Force, 1980.

[8] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach and T. Berners-Lee, "RFC 2616 - Hypertext Transfer Protocol -- HTTP/1.1," The Internet Engineering Task Force, 1999.

[9] T. Bray, J. Paoli, C. M. Sperberg-McQueen, E. Maler and F. Yergeau, "Extensible Markup Language (XML) 1.0 (Fifth Edition)," W3C, 2008.

[10] T. Berners-Lee, R. Fielding and L. Masinter, "RFC 3986 - Uniform Resource Identifier (URI): Generic Syntax," The Internet Engineering Task Force, 2005.

[11] S. Bradner, "RFC 2119 - Key words for use in RFCs to Indicate Requirement Levels," The Internet Engineering Task Force, 1997.