

TAXII 1.0 (DRAFT)

Capabilities and Services

Charles Schmidt & Mark Davidson

About This Talk

- **Look at the use scenarios we want to support and how we have designed TAXII to support them**
 - TAXII supports the sharing models people use today, but allows more automation
- **We are discussing a draft specification**
 - There are multiple open questions – we do not have all the answers
- **We want your input**
 - Please ask questions
 - Please feel free to provide suggestions for changes




What is TAXII?

- **Trusted Automated eXchange of Indicator Information**
- **The goal of TAXII is to facilitate the exchange of structured cyber threat information**
 - Specifically, TAXII is designed to support existing sharing paradigms, but do so in a more automated manner
 - “Structured cyber threat information” = STIX
- **TAXII defines the network-level activity of the exchange**
 - Defines messages to exchange data and to set up future data exchanges
 - Does NOT:
 - Dictate or control how data is handled behind the network interface
 - Dictate or control sharing policies (with whom one shares, what one shares with specific parties, etc.)
 - TAXII is NOT a sharing program

Sharing Models

- Research has identified three primary sharing models:


- Source/subscriber
- Peer-to-peer
- Hub and spoke



- TAXII can support all of these sharing models

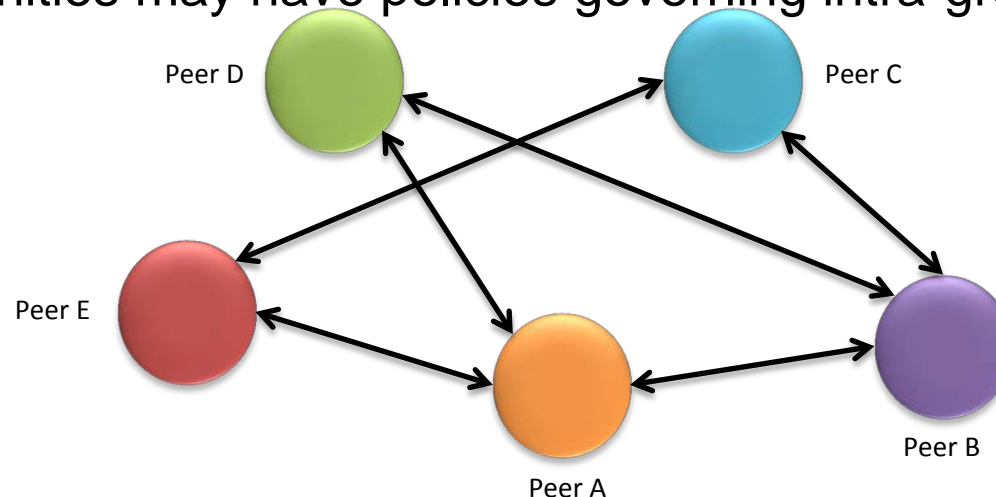
Source/Subscriber Sharing Model

- **All participants have a single role**
 - “Source” is a data producer
 - “Subscribers” are data consumers
- **Multiple distribution options**
 - “push messaging” (Analogous to subscription to mail alerts)
 - “pull messaging” (Analogous to an RSS feed)
- **Source might have multiple sharing levels**
 - Not all subscribers necessarily see the same data




Peer-to-Peer Sharing Model

- Individual participants may be data producers and/or data consumers in multiple relationships
- Multiple distribution options
 - “Push messaging” (Analogous to an email message)
 - “Pull messaging” (Analogous to a blog or similar web site)
- Individual participants decide with whom they are sharing as well as what they share
 - Communities may have policies governing intra-group sharing



Hub and Spoke Sharing Model

- **Individual participants may be data producers and/or data consumers all in a 1-1 relationship with the hub**
 - Like Source/Subscriber where subscribers also contribute data
- **Multiple distribution options**
 - “Push messaging” (Analogous to a mailing list)
 - “Pull messaging” (Analogous to a bulletin board)
- **Spokes decide what to send to the hub; Hub may make further access decisions before re-sharing with other spokes**



TAXII Services

- **TAXII defines the behavior for multiple services:**
 - Feed Management Service – receive requests for information about and for data feed subscription management
 - Inbox Service – receive pushed content
 - Poll Service – receive content pull requests
 - Discovery Service – provide information about other TAXII services
- **All TAXII services are optional – use what you need**
- **TAXII Services just dictate message exchanges**
 - Processing the details of messages is outside the scope of this specification
 - E.g., determining whether to honor a subscription request, or determining whether a piece of data should be sent to a consumer, etc.

TAXII Feed Management Service

- **Hosted by data producers**
- **Receives queries about offered TAXII data feeds**
 - Provides feed names and descriptions
 - How TAXII data feed content can be accessed (“pull” or indicate delivery protocols)
 - Any other information about a TAXII data feed (e.g., membership requirements, payment requirements, etc.)
- **Receives requests to manage TAXII data feed subscriptions**
 - Subscribe, unsubscribe, pause delivery, resume delivery, modify subscription, status query
 - TAXII does not specify the process for deciding whether to allow the requested action to occur nor how the action manifests
- **Note that the Feed Management Service does not deliver content**

TAXII Data Feeds

- **TAXII does not dictate how data producers store or organize their data...**
 - ...but **TAXII requires some common handle for communication**
- **TAXII Data Feed – a producer-dictated organization of their data**
 - A given data record might exist in one or more TAXII data feeds
 - Producers decide what data feeds represent. Examples:
 - Topic – e.g., a feed for spear-phishing, a feed for botnets, etc.
 - Subject – e.g., a feed for each identified STIX campaign
 - Access – e.g., a feed for gold-level subscribers, a feed for silver-level, etc.
 - Or producer might just have one feed with everything in it
- **In TAXII, all data distribution (push or pull) occurs relative to a TAXII Data Feed**

TAXII Inbox, Poll, and Discovery Services

■ Inbox Service

- Hosted by consumers to receive pushed content
- Basically a listener for incoming content

■ Poll Service

- Hosted by data producers
- Consumers request updates relative to a TAXII data feed
- To support this, TAXII requires all records within a TAXII data feed to be assigned a timestamp
 - Data producers can decide the meaning, if any, of the timestamp
 - Poll requests indicate a range of timestamps to collect
 - Poll responses identify returned range – recipient can track to avoid re-requesting content

■ Discovery Service

- Identify services and how to contact them

Polling vs. Querying

- **Polling allows consumers to tune requests based on data producer-declared organization of data**
 - I.e., “TAXII data feeds” and “timestamps”
- **Polling does NOT consider the contents of the data itself**
 - E.g., cannot ask for information about a specific IP address
 - Requests for records based on the record content = “querying”
- **This is a maturity issue – TAXII will support querying eventually**
 - Issue is how to usefully identify relevant STIX records

Design Principles

- **Minimize inter-session state for TAXII exchanges**
 - No exchange requires information from a previous exchange
 - TAXII back-end still needs to be stateful (e.g., record subscriptions, etc.)
- **A la carte implementation**
 - Pick the services that are useful and skip the others
- **Avoid specifying policy decision/enforcement behavior**
 - Would require standardization of policy expression – expectation was that this would be disruptive
- **Match existing procedures**
 - Follow existing sharing models
 - Minimize changes to existing infrastructure
 - TAXII does not attempt to subsume data management functions
 - Support existing technologies and mechanisms

TAXII Bindings

- **TAXII can support multiple protocols**
 - TAXII 1.0 defines the use of HTTP/HTTPS, but could define others (e.g., SMTP)
- **TAXII can support multiple data formats**
 - TAXII 1.0 defines XML bindings for messages but could define others (e.g., JSON)
- **Where appropriate, TAXII messages specify supported bindings**
 - E.g., Discovery service identifies supported protocols, etc.

Source/Subscriber Walkthrough

Background


- **One possible way to use TAXII to implement Source/Subscriber**
 - Others may make different choices

- **Assume an existing sharing arrangement**
 - A vendor (the source) publishes threat alerts as information becomes known
 - Customers (subscribers) can pay to receive these daily updates
 - Multiple levels of access depending on contract costs
 - Currently, customers log into the vendor web site to view updates
 - Customers can view the threat alerts as STIX XML documents

Step 1: Source Organizes its Data

- **Vendor organizes data records into TAXII Data Feeds**
 - Decides on “contract level” for feeds
 - Many records will be present in all feeds, but some fields may be stripped before dissemination
 - Access to a feed contingent upon the purchasing of a contract
- **Vendor labels all data within each TAXII Data Feed with a timestamp**
 - Decides to use the time of posting as that timestamp
 - More than one data record may have the same timestamp – not a problem
 - A single record could have the same timestamp in all data feeds – not a requirement

Step 2a: Source Implements TAXII Services

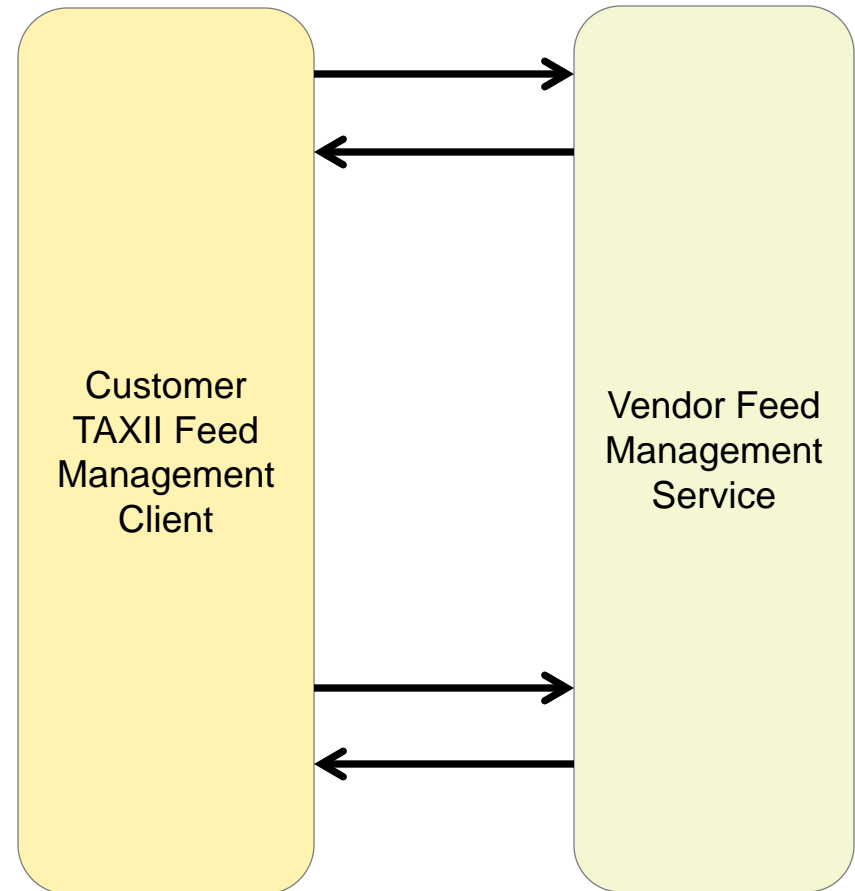
- **Decides to implement a Feed Management Service**
 - Feed Information Requests
 - Lists available feeds
 - Explain what information is provided via each feed (i.e., contract levels)
 - Reference to site where one can purchase necessary contracts
 - Feed Management Requests
 - Forward management requests to back-end for comparison to purchased contracts
 - **Decides to implement a Poll Service**
 - Give customers the option to pull content from a feed
 - **Decides to implement a TAXII Inbox Client**
 - Support pushing content to customer Inbox Services
 - **Decides NOT to implement a Discovery Service**
 - Vendor decides to continue publishing this information using HTML
- 
- MUST do at least one

Step 2b: Subscriber Implements TAXII Service

- **May implement an Inbox Service**
 - If customer wishes have updates pushed, must implement Inbox
 - Inbox listens to appropriate port for connections
 - In TAXII 1.0, this would be a (truncated) HTTP server
 - May avoid implementing if all content to be pulled via Poll Service
- **Subscribers may have a TAXII Poll Client for pull messaging**
- **For this design, subscribers must have a TAXII Feed Management Client**

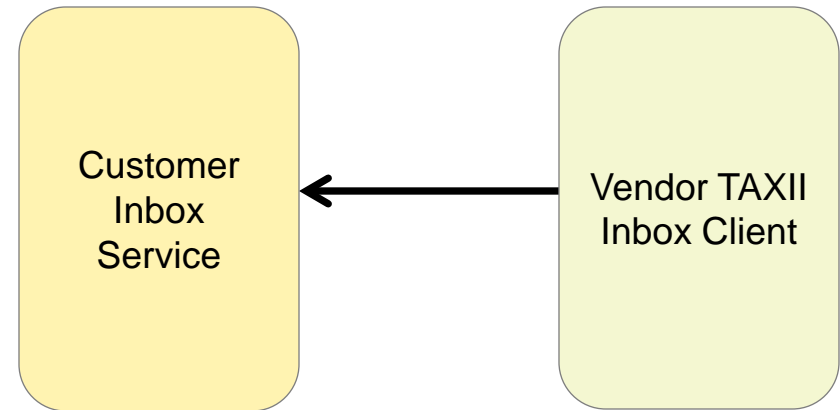
Step 3: Establish Sharing Relationships

- **Customer contacts vendor Feed Management Service to get list of feeds**
- **Customer purchases a contract via Vendor web site**
 - Also establishes authentication credentials
- **Customer contacts vendor Feed Management Service to establish subscription**
 - Request verified before acceptance

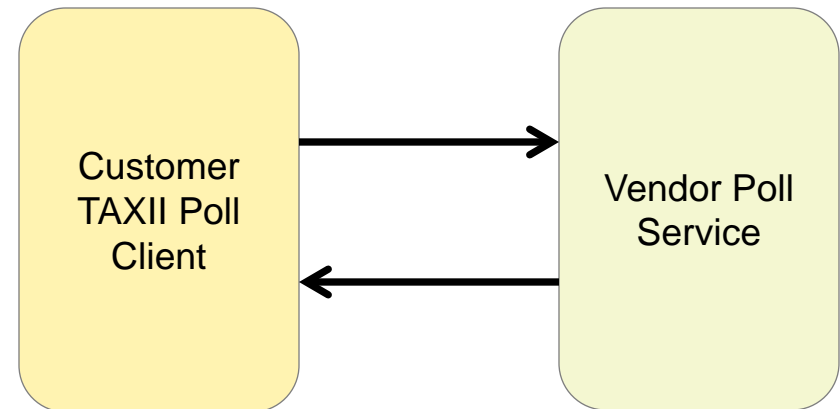


Step 4: Share

- **Content pushed to Customer's Inbox Service**



- **Customer pulls from Vendor's Poll Service**
 - Request verified before being fulfilled



Hub and Spoke Walkthrough

Background

- **One possible way to use TAXII to implement Hub and Spoke**
 - Others may make different choices

- **Assume an existing sharing arrangement**
 - Community exists with a pre-existing intra-group sharing agreement
 - Currently all threat alerts sent via e-mail to the group mailing list
 - Automatically re-distributed to all group members
 - Customers receive threat alerts as STIX XML documents in attachments

Step 1a: Hub Implements TAXII Services

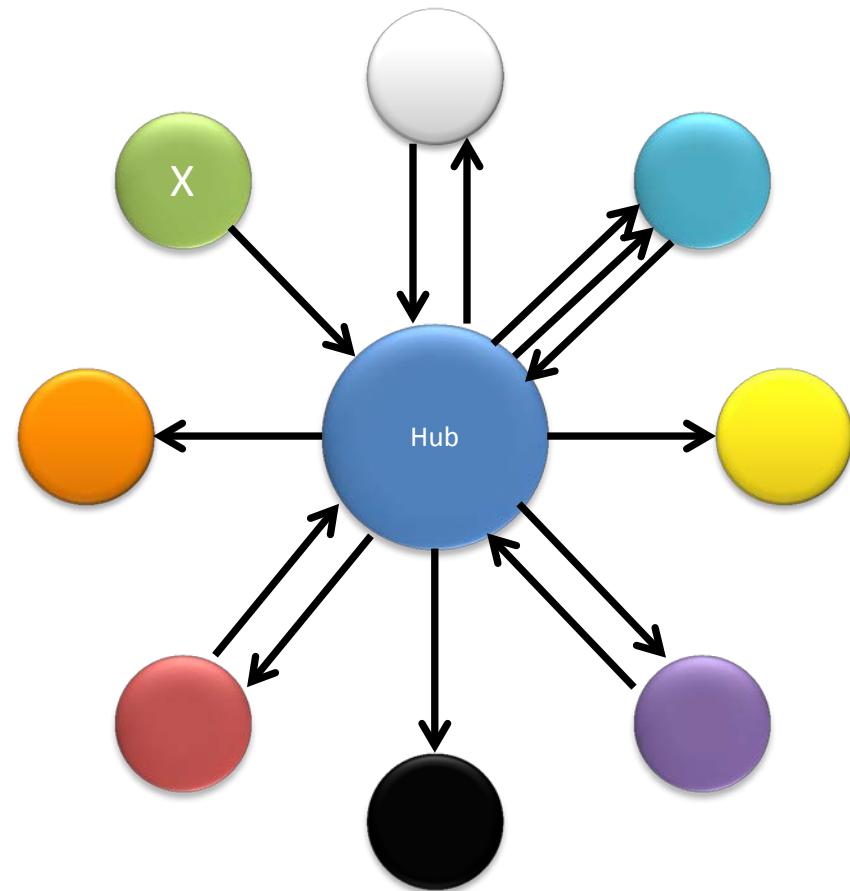
- **Decide to implement a Inbox Service**
 - Used to receive all input from spokes (Hub does not poll)
- **Decide to implement a TAXII Inbox Client for message delivery**
 - Support pushing of alerts to spokes
- **Decide to implement a Poll Service**
 - Support spokes pulling current and/or archived alerts
 - Decide on only one TAXII data feed for all information
 - Decide timestamps = the time the alert arrives in Hub's Inbox
- **Decide NOT to implement a Discovery Service**
 - Members informed of the Hub's services via other means
- **Decide NOT to implement a Feed Management Service**
 - Spokes automatically enrolled when they join the sharing group

Step 1b: Spokes Implement TAXII Services

- **Spokes that produce data implement a TAXII Inbox Client**
 - Used to send alerts to the Hub's Inbox Service
- **May implement an Inbox Service**
 - If spoke wishes have updates pushed, must implement Inbox
 - May avoid implementing if all content to be pulled via Poll Service
- **Some spokes may implement a TAXII Poll Client**
 - May avoid this use if all content to be pushed to the spoke's Inbox Service

Step 2: Share

- Spoke X pushes new alert to Hub's Inbox Service
- Hub re-sends alert to all spokes that requested push notification
- Hub archives alert so spokes can poll for the alert at a later time



What TAXII Does

- **Common behavior to automate aspects of sharing structured cyber threat information**
- **Support the primary existing sharing models**
- **Implement components as-needed**

Simplify automated sharing of structured threat information

For more information

- <http://taxii.mitre.org/>
- **Sign up for the TAXII Discussion and Announcement mailing lists**
 - <http://taxii.mitre.org/community/registration.html>
- **Related sites**
 - <https://stix.mitre.org/>
 - <http://cybox.mitre.org/>

Help out

- **TAXII 1.0 is still in DRAFT form**
- **Please tell us if TAXII is going in the right direction**
 - Does it adequately cover your use cases?
 - Are the TAXII services reasonable divisions of activity?
 - What are your thoughts on the TAXII bindings?
- **Draft specifications are available on the TAXII web site**

We need your help to make sure TAXII meets its goal of simplifying the sharing of structured threat information