THE MITRE CORPORATION

The TAXII Services Specification

Version 1.0

Mark Davidson, Charles Schmidt

04/30/2013

The Trusted Automated eXchange of Indicator Information (TAXII™) specifies mechanisms for exchanging structured cyber threat information between parties over the network. This document describes TAXII's Capabilities, Services, Messages, and Message Exchanges.

Trademark Information

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

Feedback on this or any of the other TAXII specifications is welcome and can be sent to taxii-discussion-list@lists.mitre.org after signing up on the community registration page (http://taxii.mitre.org/community/registration.html).

Comments, questions, suggestions, and concerns are all appreciated.

Table of Contents

1 Introduction

This document defines requirements for TAXII Services, TAXII Messages, and TAXII Message Exchanges. The requirements set out in this document apply to all TAXII Message Bindings and Protocol Bindings. It is recommended that the reader familiarize themself with the relationship between the TAXII specifications, as outlined in the TAXII Overview [1].

1.1 The TAXII Services Specification

This specification provides normative text on TAXII Services, Messages, and Message Exchanges. It does not provide details about how TAXII Messages are transported, leaving that to Protocol Binding Specifications. Likewise, this document identifies the information conveyed in each TAXII Message, but does not provide details about how TAXII Messages are formatted, leaving that to Message Binding Specifications.

1.1.1 TAXII Services Version ID

The TAXII Services Version ID for the version of TAXII described in this specification is:

```
urn:taxii.mitre.org:services:1.0
```

The use of this and other TAXII Version ID strings is described in Section 3.1.6.

1.2 Document Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this specification are to be interpreted as described in IETF RFC 2119. [2]

1.3 Terms and Definition

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

1.3.1 TAXII Concepts

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

Cyber Threat Information - Any information of interest to those who analyze or respond to cyber threats. This includes, but is not limited to, information about malware, threat actors, campaigns, cyber incidents, observables corresponding to a threat, and other information associated with cyber threat details.

TAXII Data Feed - A collection of structured cyber threat information that can be exchanged using TAXII. Each TAXII Data Feed has a name that uniquely identifies it among feeds from a given source of Cyber Threat Information. For more on TAXII Data Feed Names, see Section 3.1.3.

TAXII Content - A piece of structured cyber threat information. A piece of TAXII Content is considered "atomic" in that TAXII does not support sending portions of TAXII Content separately from one another.

Timestamp Label - A label in the form of a timestamp that is assigned to each piece of content within a TAXII Data Feed. For more on Timestamp Labels, see Section 3.1.5.

TAXII Message - A discrete block of information that is passed from one entity to another over the network.

TAXII Message Exchange - A defined sequence of TAXII Messages undertaken by two parties, usually in the form of a request-response pair.

TAXII Service - Functionality that is accessed or invoked through the use of one or more TAXII Message Exchanges. TAXII Services support one or more message exchanges to provide functionality.

TAXII Capability - A high-level activity supported by TAXII through the use of one or more TAXII Services.

1.3.2 TAXII Roles

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

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

Consumer - An entity that is the recipient of structured cyber threat information.

Note that these roles are not mutually exclusive - one entity might be both a Consumer and a Producer of structured cyber threat information.

1.3.3 TAXII Network Components

The following terms are used to define the components of a TAXII Implementation using a typical client-server model.

TAXII Implementation - A specific implementation of a TAXII Architecture.

TAXII Daemon - The part of a TAXII Implementation that provides one or more TAXII Services. To support this functionality, it is assumed that a TAXII Daemon is persistently listening for new TAXII requests over a network.

TAXII Client - The part of a TAXII Implementation that initiates an exchange with a remote TAXII Daemon. 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 Daemon and disconnect from the network when this interaction has concluded.

Note that TAXII Network Components do not map directly to the TAXII Roles previously defined: For example, an entity might both host a TAXII Daemon 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.

2 TAXII Services

TAXII Services represent a set of mechanisms necessary to support some TAXII Capability or Capabilities. A TAXII Implementation can support some, all, or even none of the defined TAXII Services. (On the latter note, one can still make use of some TAXII Capabilities without ever hosting a TAXII Daemon that supports any of the described TAXII Services.)

This section defines the following Services:

- Discovery Service – Provide information about offered TAXII Services.
- Feed Management Service – Support management of TAXII Data Feed subscriptions.
- Inbox Service – Support Producer-initiated pushes of cyber threat information.
- Poll Service - Support Consumer-initiated pulls of cyber threat information.

The following sections look at each of these services in more detail.

2.1 Discovery Service

The Discovery Service is the mechanism for communicating information related to the availability and use of TAXII Services. The Discovery Service provides a requester with a list of TAXII Services and how these Services can be invoked (i.e., the address of the TAXII Daemon that implements that service and the bindings that Daemon supports). A single Discovery Service might report on TAXII Services hosted by TAXII Daemons on multiple endpoints or even across multiple organizations - the owner of a Discovery Service can define its scope as they wish, as long as they comply with legal, ethical, and other considerations. A Discovery Service is not required to disclose all TAXII Services of which it is aware; a Discovery Service can use a variety of factors to determine which Services to disclose to the requester, including but not limited to the requester's identity. In order to facilitate automation, each TAXII Protocol Binding Specification defines a recommended default address for the Discovery Service.

A Discovery Service implementation MUST support the Discovery Exchange as defined in Section 4.1.2.

2.2 Feed Management Service

The Feed Management Service is the mechanism by which a Consumer can request information about TAXII Data Feeds, request subscriptions to TAXII Data Feeds, request the status of a subscription, or terminate existing subscriptions to TAXII Data Feeds. The Feed Management Service does not deliver TAXII Data Feed content (i.e., the threat information the Producer publishes in association with the named TAXII Data Feed). Instead, TAXII Data Feed content is either sent to a Consumer's TAXII Daemon implementing an Inbox Service in Producer-initiated exchanges or in direct response to Consumer requests to the Producer's Poll Service.

A Feed Management Service implementation MUST support at least one of the Feed Information Exchange or the Subscription Management Exchange, as defined in Sections 4.1.3 and 4.1.4, respectively.

A Feed Management Service implementation MAY support both the Feed Information Exchange and the Subscription Management Exchange.

2.3 Inbox Service

The Inbox Service is the mechanism by which a Consumer accepts messages from a Producer in Producer-initiated exchanges (i.e., push messaging). A Consumer can implement this Service in order to receive TAXII Data Feed content via Producer-initiated exchanges. Such content might be the result of the Consumer's establishment of subscriptions on a Producer or can be unsolicited data.

An Inbox Service implementation MUST support the Inbox Exchange, as defined in Section 4.1.1.

2.4 Poll Service

The Poll Service is the mechanism by which a Producer allows Consumer-initiated pulls from a TAXII Data Feed (i.e., pull messaging). A Consumer contacts the Poll Service to explicitly request TAXII Data Feed content. Consumers can contact the Poll Service to request TAXII Data Feed content at the Consumer's convenience. Note that Producers can choose to offer TAXII Data Feed content through a combination of Producer-initiated pushes to the Consumer's Inbox Service and Consumer-initiated pulls from the Producer's Poll Service.

A Poll Service implementation MUST support the Feed Poll Exchange, as defined in Section 4.1.5.

3 TAXII Messages

This section defines TAXII Messages, their contents and their purposes. Some messages, such as the TAXII Status Message, are used in multiple message exchanges while others are only used in a single message exchange. The messages defined here are the only allowed messages that can be sent as part of a TAXII message exchange. While the values of some fields can be customized by implementers, implementers MUST NOT create new message types.

3.1 Message Concepts

This section contains requirements and information for concepts applicable to all TAXII Messages.

3.1.1 Message IDs

Every TAXII Message has a Message ID field. Message IDs are used to link requests with responses. Specifically, if TAXII Message B is sent in response to TAXII Message A, Message B MUST contain an "In Response To" field whose value is the Message ID of Message A. This allows the recipient of Message B to know to which of their requests it is a response.

A message sender MUST NOT reuse a particular Message ID if it is still expecting a response to an earlier request that used that same Message ID as this could lead to confusion as to which message a given response was responding to.

Message IDs MUST be an unsigned integer.

3.1.2 Extended Headers

All TAXII Messages support the use of extended headers to allow extensions to TAXII Messages. Extended headers consist of name-value pairs. Names of extended headers MUST conform to URI formatting rules [3]. In order to avoid accidental name collisions, extended header names SHOULD contain an "authority" part that identifies the entity that controls the meaning of this extended header.

Values for extended headers are unrestricted and can contain any characters and can even contain structured content. Note that some TAXII Message Bindings might prohibit certain characters or require that certain characters be escaped before the value is encoded in a TAXII Message. Individual TAXII Message Bindings indicate such requirements.

3.1.3 Data Feed Names

Every TAXII Data Feed has a unique identifier relative to the other TAXII Data Feeds from the same Producer. Different Producers can use the same Feed Name unless those Producers share a Feed Management or Poll Service.

Consumers use Feed Names as handles to a Producer's TAXII Data Feeds in their request messages. Note that because Feed Names are unique relative to a Producer rather than globally unique, it is possible that a single Consumer might interact with multiple Producers and, during the course of these interactions, encounter two distinct TAXII Data Feeds with identical Feed Names. For this reason, Consumers need to track both the Feed Name and the associated Producer identity together since the combination of these values is globally unique.

Data Feed names MUST conform to URI formatting rules [3].

3.1.4 Subscription ID

TAXII Consumers can establish subscriptions to TAXII Data Feeds provided by TAXII Producers. For convenience when manipulating existing subscriptions, TAXII defines Subscription IDs. When a Consumer successfully establishes a subscription on a Producer, the Producer assigns that subscription a Subscription ID value. From then on, both the Consumer and Producer refer to this subscription in messages using this Subscription ID value. Two subscriptions to the same TAXII Data Feed by the same Consumer MUST NOT be given the same Subscription ID.

Subscription IDs MUST conform to URI formatting rules [3].

3.1.5 Timestamp Labels

Timestamp Labels are used to give an ordering to the content within a TAXII Data Feed. Each piece of content within a TAXII Data Feed is assigned a Timestamp Label. Multiple pieces of content MUST NOT be assigned the same Timestamp Label unless they are added to the associated TAXII Data Feed as an "atomic" action. (This is necessary to prevent a race condition where a requester receives some of the content associated with a single Timestamp Label but not other content with that Timestamp Label because the request arrived part-way through the addition of this set of content.) While a Timestamp Label is in the form of a timestamp, it is important to note that Timestamp Labels do not necessarily correspond to any chronological event nor do they necessarily align with timestamps that appear within

the content of a TAXII Data Feed. The Timestamp Label is just a label, rather than a reference to some meaningful chronological time.

Timestamp Labels MUST conform to a specific set of rules:

1. Timestamp Labels MUST comply with the date-time construct as defined in IETF RFC 3339 [4].
2. Each piece of content in a TAXII Data Feed MUST have a Timestamp Label.
3. When a new piece of content (or set of content) is added to a TAXII Data Feed, that content MUST be assigned a Timestamp Label later than the Timestamp Label of any other piece of content within that feed. Note that this property MUST be maintained even if the Producer assigns Timestamp Labels that use different time zones: new Timestamp Labels MUST be chronologically later than all other previous Timestamp Labels within that TAXII Data Feed. (In other words, one can use Timestamp Labels to create an unambiguous ordering of content within a TAXII Data Feed.)
4. A Timestamp Label MAY have between 0 and 6 decimal places in its fractional second. A Timestamp Label MUST NOT contain more than 6 decimal places in its fractional second.

3.1.6 Message Version IDs, Protocol Version IDs, and Content Binding IDs

This document makes references to TAXII "Version IDs", specifically TAXII Services Version IDs, TAXII Protocol Binding Version IDs, and TAXII Message Binding Version IDs. The TAXII Version IDs are used in certain TAXII Message fields to denote specific versions of TAXII specifications. Each TAXII specification identifies its own TAXII Version ID. Different versions of each specification provide a different Version ID. Version IDs can be referenced in TAXII Message fields as a way to identify specific versions of TAXII and its bindings.

Similarly, a Content Binding ID is a string that identifies the format and version of a piece of content in a TAXII Message. The TAXII Content Binding Reference [5] defines a canonical list of Content Binding IDs for a core set of supported message content types. Content Binding IDs can be referenced in TAXII Message fields as a way to identify specific types of content in TAXII Messages.

In addition to the Version IDs defined in TAXII Specifications and the Content Binding IDs defined in the TAXII Content Binding Reference, third parties can define their own Message Binding Version IDs, Protocol Binding Version IDs, and Content Binding IDs. Third parties MUST NOT define alternate TAXII Services Version IDs. These are used to indicate custom message bindings, custom protocol bindings, or types of contents not covered in the TAXII Content Binding Reference. Third party IDs MUST conform to URI formatting rules [3]. Such URIs MUST contain an authority component (i.e., a domain name) indicating the authority that controls the meaning of the ID.

3.2 TAXII Message Representation Conventions

This section is limited to a description of the data models representing TAXII Messages. This section does not prescribe any particular binding for this data model - such details are provided by TAXII Message Binding Specifications. This section describes what information a TAXII Message conveys, while the TAXII Message Binding Specifications define how to express that information. As a result, there are not always

one-to-one mappings between fields in the data model and fields in the data bindings. For example, some bindings might require multiple field structures (e.g., elements and attributes in an XML [6] binding) to account for the intended meaning of a single field as described in this document. Alternatively, a field's value might be conveyed without any transmitted structure. For example, an XML binding might specify default values for some field structures allowing that field to be dropped from the actual message structure during communications. It is important to keep in mind that this section describes the conceptual fields in the data model; the message bindings follow those concepts, but might include structural differences to account for limitations or capabilities of the particular binding. Implementers need to consult the appropriate TAXII Message Binding Specification for binding requirements and details.

All TAXII Messages consist of two parts: a header and a body. The header contains information relevant to all message body types, while the body contains information relevant to a particular message type. The following sections describe the use of the header and body types and list their fields. Each field is listed with the following information:

- Name - A handle by which the TAXII specifications refer to this field. This might not be exactly identical to the structural field names (e.g., XML element or attribute names) that appear in the TAXII Message Binding Specifications.
- Required? - Whether the message MUST convey the indicated information. Note that in a particular message binding default values might allow a required field to be absent in the actual exchanged content. The fact that the default value is implicitly conveyed fulfills the requirement for the field's presence.
- Multiple? - Whether field indicates a single value or whether it can indicate multiple values.
- Description - A description of the information the field conveys between the message sender and recipient.

Details such as the data type of the field and the definition of controlled vocabularies used by a field are outside the scope of this document and are instead covered in each TAXII Message Binding Specification. Some fields are noted as having "sub-fields" - this is simply an organizational convenience for this document and not a requirement imposed on their representation in any given binding. The "Required?" and "Multiple?" values for a given sub-field reflect its use only within its parent field. A sub-field might not allow multiple values, but the sub-field is still able to appear and hold a single value in each of the multiple instances of its parent field.

3.3 TAXII Header

This section defines the conceptual model for the header fields of a TAXII Message.

Table 1 - TAXII Header Fields

Name	Required?	Multiple?	Description
Message ID	Yes	No	A value identifying this message.

Name	Required?	Multiple?	Description
Message Body Type	Yes	No	The type of the TAXII Message. Only identifiers for defined TAXII Messages, as defined in Section 3.4, are allowed in this field. (I.e., third parties MUST NOT define their own TAXII Message Body Types.)
In Response To	Yes, if this message is a response.	No	Contains the Message ID of the message to which this is a response, if applicable.
Extended-Header	No	Yes	Third parties MAY define their own additional header fields. Extended-Header fields that are not recognized by a recipient SHOULD be ignored. Requirements for Extended-Header fields are listed in Section 3.1.2.
Signature	No	No	This field contains a cryptographic signature for this TAXII Message. The scope of this signature is the entire TAXII Message (i.e., Signatures contained in this field can sign all or any parts of the TAXII Message). Details for how a signature is expressed are covered in each TAXII Message Binding Specification.

3.4 TAXII Message Bodies

TAXII Message bodies are used to support specific TAXII Message Exchanges. The Message Body Types defined in this specification are:

- TAXII Status Message
- TAXII Discovery Request
- TAXII Discovery Response
- TAXII Feed Information Request
- TAXII Feed Information Response
- TAXII Manage Feed Subscription Request
- TAXII Manage Feed Subscription Response
- TAXII Poll Request
- TAXII Poll Response
- TAXII Inbox Message

Each permissible TAXII Message Body Type is described in detail in the following sub-sections.

3.4.1 TAXII Status Message

A TAXII Status Message is used to indicate a condition of success or error. Status Messages are always sent from a TAXII Daemon to a TAXII Client in response to a TAXII Message. A TAXII Status Message can be used to indicate an error when to performing some requested action. It is also used in the Inbox Exchange to indicate successful reception of an Inbox Message. (See Section 4.1.1 for more on the Inbox Exchange.) Error conditions can occur because the request itself was invalid or because the recipient was unwilling or unable to honor the request.

Table 2 - TAXII Status Message Fields

Name	Required?	Multiple?	Description
Status Type	Yes	No	One of the Status Types defined in Table 3 or a third party-defined Status Type.
Status Detail	Per Status Type	No	A field for additional information about this status in a machine-readable format. (The details of this format appear in the appropriate TAXII Message Binding Specification.) The individual Status Types indicate the appropriate values for this field (if anything). For Status Types defined in Table 3, this field SHOULD only be present when the Status Type indicates and MUST only contain the indicated information. For third party defined Status Types, a Status Detail message MAY be defined.
Message	No	No	Additional information for the status. There is no expectation that this field be interpretable by a machine; it is instead targeted to a human operator.

TAXII Daemons reporting an error condition SHOULD provide as much detail as possible in the Message field. Table 3 provides canonical Status Types for TAXII Status Messages. The description of each field indicates whether a Status Detail value is defined for that Status Type. For Status Types for which Status Detail values are suggested, senders MUST NOT provide different or additional values in the Status Detail field. Each TAXII Message Binding Specification provides formatting details for each of the suggested Status Detail values.

Table 3 - TAXII Status Types

Status Type	Description
Success	The message sent was interpreted by the TAXII Daemon and the requested action was completed successfully. Note that some request messages have a corresponding response message used to indicate successful completion of a request. In these cases, that response message MUST be used instead of sending a Status Message of type "Success".
Bad Message	The message sent could not be interpreted by the TAXII Daemon (e.g., it was malformed and could not be parsed).
Denied	This is used in cases where the TAXII Client's action is being denied for reasons other than a failure to provide appropriate authentication credentials. For example, a Feed Management Service might limit the number of subscriptions a given Consumer is allowed to create. In this case, if a Consumer attempts to create a too many subscriptions, a TAXII Daemon might send a Status Message of type "Denied".
Failure	A general indication of failure. This might indicate some problem that does not have a defined Status Type, but MAY also be sent in place of any other TAXII Status Messages if a TAXII Daemon does not wish to disclose details for the failure of a request.
Not Found	The request named a target (e.g., a TAXII Data Feed name) that does not exist on the TAXII Daemon.

Status Type	Description
Polling Not Supported	The requester attempted to create a subscription where the requester polls for content, but the associated TAXII Data Feed is not available via polling. (I.e., the TAXII Data Feed is not hosted by a Poll Service.)
Retry	The request cannot be completed immediately. The Error Detail field SHOULD contain a timestamp indicating when a retry of the request might be successful. The requested action will not occur until and unless the request is repeated.
Unauthorized	The requested activity requires authentication, but either the TAXII Client did not provide authentication or their authenticated identity did not have appropriate access rights. (Note that any authentication credentials are provided at the protocol level rather than as part of a TAXII Message.)
Unsupported Message Binding	The requester identified a set of message bindings to be used in the fulfillment of its request, but none of those message bindings are supported for the requested action. The Status Detail field SHOULD contain a list of acceptable message bindings.
Unsupported Content Binding	The requester identified a set of content bindings to be used in the fulfillment of its request, but none of those content bindings are supported for the requested action. The Status Detail field SHOULD contain a list of acceptable Content Binding IDs.
Unsupported Protocol Binding	The requester identified a set of protocol bindings to be used in the fulfillment of its request, but none of those protocol bindings are supported for the requested action. The Status Detail field SHOULD contain a list of acceptable protocol bindings.

3.4.1.1 Third Party Status Types

Third parties MAY define additional Status Types to indicate error conditions instead of using one of the defined Status Types provided in Table 3. Third party Status Types can be used to indicate an error condition specific to a particular TAXII implementation or user group. If the recipient does not recognize a third party Status Type, it SHOULD be treated as a Status Type of "Failure". For this reason, third parties MUST NOT define additional Status Types to indicate success conditions.

Status Types defined by a third party MUST conform to URI formatting rules [3]. In order to avoid accidental name collisions, third party defined Status Types SHOULD contain an "authority" part that identifies the entity that controls the meaning of this Status Type. Third parties MUST NOT redefine the meaning of the canonical Status Types provided in Table 3.

Status Types defined by a third party MAY make use of the Status Detail field to provide machine readable information about the given status condition. The party defining the new Status Type is responsible for determining the nature of appropriate Status Detail information, and all users that utilize this new status type SHOULD conform to those guidelines. It is recommended that Status Detail contents be expressible as simple strings to improve compatibility across different TAXII Message Bindings.

3.4.2 TAXII Discovery Request

This message is sent to a Discovery Service to request information about provided TAXII Services. Such information includes what TAXII Services are offered, how the TAXII Daemons that support those Services can be accessed, and what protocols and message bindings are supported. The body of this message is empty.

3.4.3 TAXII Discovery Response

This message is sent from a Discovery Service in response to a TAXII Discovery Request if the request is successful. If there is an error condition, a TAXII Status Message indicating the nature of the error is sent instead.

Table 4 - TAXII Discovery Response Message Fields

Name	Required?	Multiple?	Description
Service Instance	No	Yes	This field MAY appear any number of times (including 0), each time identifying a different instance of a TAXII Service. This field has several sub-fields.
Service Type	Yes	No	This field identifies the Service Type of this Service Instance (e.g., Poll, Inbox, Feed Management, or Discovery).
Services Version	Yes	No	This field identifies the TAXII Services Specification to which this Service conforms. This MUST be a TAXII Services Version ID as defined in a TAXII Services Specification.
Protocol Binding	Yes	No	This field identifies the protocol binding supported by this Service. This MUST be a TAXII Protocol Binding Version ID as defined in a TAXII Protocol Binding Specification or by a third party.
Service Address	Yes	No	This field identifies the network address of the TAXII Daemon that hosts this Service. The Service Address MUST use a format appropriate to the Protocol Binding field value.
Message Binding	Yes	Yes	This field identifies the message bindings supported by this Service instance. Each message binding MUST be a TAXII Message Binding Version ID as defined in a TAXII Message Binding Specification or by a third party.
Inbox Service Accepted Content	Only if the Service Type is an Inbox Service	Yes	This field SHOULD NOT be present for any Service Type other than Inbox; recipients MUST ignore this field if the Service Type is not Inbox, as it has no meaning. This field identifies content bindings that this Inbox Service is willing to accept. Each Inbox Service Accepted Content MUST be a Content Binding ID as defined in the TAXII Content Binding Reference or by a third party. A special value is provided by all TAXII Message Binding Specifications to indicate that the Inbox Service accepts content using any content binding.
Available	No	No	This field indicates whether the identity of the requester (authenticated or otherwise) is allowed to access this TAXII Service. It can indicate that the requester is known to have access, known not to have access, or that access is unknown.

Name		Required?	Multiple?	Description
	Message	No	No	This field contains a message regarding this Service instance. This message is not required to be machine readable and is usually a message for a human operator.

Each Service Instance record identifies one instance of a TAXII Service as hosted by a particular TAXII Daemon. Each instance identifies a single TAXII Protocol Binding Specification and a network address for that binding. An instance MAY identify multiple TAXII Message Binding Specifications and (if the TAXII Service is an Inbox Service) a set of content bindings. Note that, within a single Service Instance record, it is expected that every combination of message bindings and content bindings is acceptable. In other words, if the record for an Inbox Service lists two acceptable message bindings (1 and 2) and three acceptable content bindings (A, B, and C), all six message binding-content binding combinations are considered supported (1A, 1B, 1C, 2A, 2B, and 2C).

If a given Inbox Service only accepts certain combinations of message bindings and content bindings, multiple Service Instance records can be created for this one service to avoid incorrectly indicating support for an unsupported combination. For example, if a particular Inbox Service supported two message bindings (1 and 2) and three content bindings (A, B, and C), but only supported the a subset of all possible combinations (e.g., it only supported 1A, 1B, 2B, and 2C), this service would need to be represented by multiple Service Instance records (i.e., one record that noted support for message binding 1 and content bindings A and B, and a second record that noted support for message binding 2 and content bindings B and C). This situation only arises when an instance of an Inbox Service supports multiple message and content bindings but fails to support all combinations of the two.

Note that the Discovery Service is not required to list all existing TAXII Services of which it is aware. For example, some TAXII Services might only be publicized to specific, authenticated parties. As such, different requesters might get different responses to a Discovery Request sent to the same Discovery Service.

3.4.4 TAXII Feed Information Request

This message is sent to a Feed Management Service to request information about the available feeds. The body of this message is empty.

3.4.5 TAXII Feed Information Response

This message is sent in response to a TAXII Feed Information Request if the request is successful. If there is an error condition, a TAXII Status Message indicating the nature of the error is sent instead. Note that the Producer is under no obligation to list all feeds and can exclude any or all feeds from this response for any reason. For example, the Producer might wish to exclude feeds created for a specific customer from a list of all feeds. As such, different requesters might be given different lists of feeds to their requests to the same Feed Management Service.

Table 5 - TAXII Feed Information Response Fields

Name	Required?	Multiple?	Description
Feed Information	No	Yes	This field MAY appear any number of times (including 0), each time identifying a different TAXII Data Feed. It has several sub-fields.
Feed Name	Yes	No	This field contains the name by which this TAXII Data Feed is identified. Each TAXII Data Feed managed by a single Feed Management Service MUST have a unique Feed Name.
Feed Description	Yes	No	This field contains a prose description of this TAXII Data Feed. This field can explain how to gain access to this TAXII Data Feed if access is restricted. (E.g., pay a fee, only available to members of some organization, etc.)
Supported Content	Yes	Yes	This field contains Content Binding IDs indicating which types of content are currently expressed in this TAXII Data Feed. Each Supported Content MUST be a Content Binding ID as defined in the TAXII Content Binding Reference or by a third party.
Available	No	No	This field indicates whether the identity of the requester (authenticated or otherwise) is allowed to access this TAXII Service. It can indicate that the requester is known to have access, known not to have access, or that access is unknown.
Push Method	Yes if there are no instances of the Polling Service Instance field.	Yes	This field identifies the protocols that can be used to push content via a subscription. This field MAY appear multiple times if content from this TAXII Data Feed can be pushed via multiple protocols. This field has multiple sub-fields. At least one instance of this field or the Poll Service Instance field MUST be present. Both MAY be present.
Push Protocol	Yes	No	This field identifies a protocol binding that can be used to push content to an Inbox Service instance. This MUST be a TAXII Protocol Binding Version ID as defined in a TAXII Protocol Binding Specification or by a third party.
Push Message Binding	Yes	Yes	This field identifies the message bindings that can be used to push content to an Inbox Service instance using the protocol identified in the Push Protocol field. Each message binding MUST be a TAXII Message Binding Version ID as defined in a TAXII Message Binding Specification or by a third party.

Name	Required?	Multiple?	Description
Polling Service Instance	Yes if there are no instances of the Push Method field.	Yes	This field identifies the bindings and address a Consumer can use to interact with a Poll Service instance that supports this TAXII Data Feed. This field MAY appear multiple times if multiple Poll Services support this TAXII Data Feed. A subscription might or might not be required before content from this data feed can be polled. This field has multiple sub-fields. At least one instance of this field or the Push Method field MUST be present. Both MAY be present.
Poll Protocol	Yes	No	This field identifies the protocol binding supported by this Poll Service instance. This MUST be a TAXII Protocol Binding Version ID as defined in a TAXII Protocol Binding Specification or by a third party.
Poll Address	Yes	No	This field identifies the address of the TAXII Daemon hosting this Poll Service instance. This field MUST use a format appropriate to the Poll Protocol field value.
Poll Message Binding	Yes	Yes	This field identifies the message bindings supported by this Poll Service instance. Each message binding MUST be a TAXII Message Binding Version ID as defined in a TAXII Message Binding Specification or by a third party.
Subscription Method	No	Yes	This field identifies the protocol and address of the TAXII Daemon hosting the Feed Management Service that can process subscriptions for this TAXII Data Feed. This field MUST be absent if there is not a TAXII Service that processes subscription requests for this feed. In that case subscriptions, if supported, would need to be established by mechanisms other than TAXII. In the case of alternative subscription methods, the Feed Description field could provide procedures for subscribing.
Subscription Protocol	Yes	No	This field identifies the protocol binding supported by this Feed Management Service instance. This MUST be a TAXII Protocol Binding Version ID as defined in a TAXII Protocol Binding Specification or by a third party.
Subscription Address	Yes	No	This field identifies the address of the TAXII Daemon hosting this Feed Management Service instance. This field MUST use a format appropriate to the Subscription Protocol field value.
Subscription Message Binding	Yes	Yes	This field identifies the message bindings supported by this Feed Management Service Instance. Each message binding MUST be a TAXII Message Binding Version ID as defined in a TAXII Message Binding Specification or by a third party.

3.4.6 TAXII Manage Feed Subscription Request

This message is used to manage (i.e., subscribe, unsubscribe, or request the status of) a subscription. The Feed Management Service responds with a TAXII Manage Feed Subscription Response if the request is successful and will be honored or with a TAXII Status Message if the request is being rejected or there was an error.

Table 6 - TAXII Manage Feed Subscription Request Fields

Name	Required?	Multiple?	Description
Feed Name	Yes	No	This field identifies the name of the TAXII Data Feed to which the action applies. Each TAXII Data Feed managed by a single Feed Management Service MUST have a unique Feed Name.
Action	Yes	No	This field identifies the requested action to take. The action MUST be one of the following: ○ SUBSCRIBE - Request a subscription to the named TAXII Data Feed. ○ UNSUBSCRIBE - Request cancellation of an existing subscription to the named TAXII Data Feed. ○ STATUS - Request information on all subscriptions the requester has established for the named TAXII Data Feed. No subscription state is changed in response to this action.
Subscription ID	Per Action	No	This field contains the ID of a previously created subscription. For the UNSUBSCRIBE action this field MUST be present. This field MUST be ignored if present in a SUBSCRIBE or STATUS action message.
Delivery Parameters	Yes	No	This field identifies the delivery parameters for this request. This field contains multiple sub-fields. If the subscription action is SUBSCRIBE, subfields indicates how the requester is requesting to have messages pushed to their Inbox Service. A special value is provided by all TAXII Message Binding Specifications to indicate that the requester is not requesting pushed content and will poll for subscription content instead use a Poll Service hosted by the data provider. In this case, if the TAXII Data Feed cannot be polled, a Status Message with a status of 'Polling Not Supported' SHOULD be returned. For actions of UNSUBSCRIBE and STATUS, this field MUST be ignored by recipients and SHOULD NOT be included by senders.
Inbox Protocol	Yes, if requesting push messaging	No	This field identifies the protocol to be used when pushing TAXII Data Feed content to a Consumer's TAXII Inbox Service implementation. If the Data Feed does not support the named Inbox Protocol, a Status Message with a status of 'Unsupported Protocol Binding' SHOULD be returned. The Inbox Protocol MUST be a TAXII Protocol Binding Version ID as defined in a TAXII Protocol Binding Specification or by a third party.

Name	Required?	Multiple?	Description
Inbox Address	Yes, if requesting push messaging	No	This field identifies the address of the TAXII Daemon hosting the Inbox Service to which the Consumer requests content for this TAXII Data Feed to be delivered. The address MUST be of the appropriate type for the network protocol identified in the Inbox Protocol field.
Delivery Message Binding	Yes, if requesting push messaging	No	This field identifies the message binding to be used to send pushed content for this subscription. If the TAXII Data Feed does not support the Delivery Message Binding, a Status Message with a status of 'Unsupported Message Binding' SHOULD be returned. Each Delivery Message Binding MUST be a TAXII Message Binding Version ID as defined in a TAXII Message Binding Specification or by a third party.
Content Binding	Yes, if requesting push messaging	Yes	This field contains Content Binding IDs indicating which types of contents the Consumer requests to receive for this TAXII Data Feed. This field MUST contain a Content Binding IDs as defined in the TAXII Content Binding Reference or by a third party. If none of the listed Content Binding values are supported by the Data Feed, a Status Message with a status of 'Unsupported Content Binding' SHOULD be returned. A special value is provided by all TAXII Message Binding Specifications to indicate that all content bindings are accepted.

Responses to subscription management requests MUST be processed using the following criteria in order:

1. Any attempt to manage subscriptions that require authentication where the request comes from a source that lacks appropriate authentication SHOULD result in an appropriate TAXII Status Message (normally "Unauthorized") without changing existing subscriptions. This takes precedence over all other conditions.
2. Attempts to manage feeds where the requested Feed Name does not correspond to an existing Feed Name SHOULD result in an appropriate TAXII Status Message (normally "Not Found") without changing existing subscriptions.
3. Attempts to unsubscribe (UNSUBSCRIBE action) where the Subscription ID does not correspond to an existing subscription on the named TAXII Data Feed by the identified Consumer SHOULD be treated as a successful attempt to unsubscribe and result in a TAXII Manage Feed Subscription Response without changing existing subscriptions. In other words, the requester is informed that there is now no subscription with that Subscription ID (even though there never was one in the first place).
4. Attempts to create a new subscription (SUBSCRIBE action) where the requested Inbox Protocol, Delivery Message Binding, or Content Binding of the subscription to be created is not supported SHOULD result in an appropriate TAXII Status Message (normally "Unsupported Protocol

Binding", "Unsupported Message Binding", or "Unsupported Content Binding" respectively) without changing existing subscriptions.

5. Attempts to create a new subscription (SUBSCRIBE action) where the subscription to be created is identical to an existing subscription (i.e., same Feed Name and Delivery Parameters) SHOULD result in a TAXII Manage Feed Subscription Response that returns that existing subscription's Subscription ID without changing existing subscriptions. That is, the Feed Management Service SHOULD NOT create exact duplicates of existing subscriptions, but the client SHOULD be informed that the requested subscription is established.

3.4.7 TAXII Manage Feed Subscription Response

This message is returned in response to a TAXII Manage Feed Request Message if the requested action was successfully completed.

Table 7 - TAXII Manage Feed Subscription Response Fields

Name	Required?	Multiple?	Description
Feed Name	Yes	No	This field identifies the name of the TAXII Data Feed to which the action applies. Each TAXII Data Feed managed by a single Feed Management Service MUST have a unique Feed Name.
Message	No	No	This field contains a message associated with the subscription response. This message is not required to be machine readable and is usually a message for a human operator.
Subscription Instance	Yes	Yes	This field contains information about existing subscriptions by the requester to the given TAXII Data Feed. It appears any number of times (including 0) if this message is in response to a STATUS action, or exactly once if responding to any other action.
Subscription ID	Yes	No	This field contains an identifier that is used to reference the given subscription in subsequent exchanges.
Delivery Parameters	Yes if message is responding to a STATUS action	No	This field contains a copy of the Delivery Parameters of the Manage Feed Subscription Request Message that established this subscription. This field MUST be present if this message is responding to a STATUS action. This field MUST NOT be present when responding to other actions.
Inbox Protocol	Yes, if requested push messaging	No	This field contains a copy of the Inbox Protocol field in the Manage Feed Subscription Request Message that established this subscription.
Inbox Address	Yes, if requested push messaging	No	This field contains a copy of the Inbox Address field in the Manage Feed Subscription Request Message that established this subscription.

Name	Required?	Multiple?	Description
Delivery Message Binding	Yes, if requested push messaging	No	This field contains a copy of the Delivery Message Binding field in the Manage Feed Subscription Request Message that established this subscription.
Content Binding	Yes, if requested push messaging	Yes	This field contains a copy of the Content Binding field(s) in the Manage Feed Subscription Request Message that established this subscription.
Poll Instance	Yes, if action was SUBSCRIBE and the request was for polling. Optional otherwise	Yes	Each Poll Instance represents an instance of a Poll Service that can be contacted to retrieve content associated with the new Subscription. If the Manage Feed Subscription Request Message indicated that the requester wished to poll for content there MUST be at least one Poll Instance in the response to a SUBSCRIBE action. This field indicates where Poll Request Messages can be sent for the given subscription. If the requester is requesting pushed content the Subscription Response Message MAY contain one or more Poll Instances if the subscriber is also allowed to poll for content in addition to receiving pushed content. This field MAY be present for requests actions other than SUBSCRIBE.
Poll Protocol	Yes	No	The protocol binding supported by this instance of a Polling Service. This field MUST be a TAXII Protocol Binding Version ID as defined in a TAXII Protocol Binding Specification or by third parties.
Poll Address	Yes	No	This field identifies the address of the TAXII Daemon hosting this Poll Service. This field MUST use a format appropriate to the Poll Protocol field value.
Poll Message Binding	Yes	Yes	This field identifies one or more message bindings that can be used when interacting with this Poll Service instance. Each message binding MUST be a TAXII Message Binding Version ID as defined in a TAXII Message Binding Specification or by a third party.

3.4.8 TAXII Poll Request

This message is sent from a Consumer to a TAXII Poll Service to request that data from the TAXII Data Feed be returned to the Consumer. Poll Requests are always made against a specific TAXII Data Feed. Whether or not the Consumer needs an established subscription to that TAXII Data Feed in order to receive content is left to the Producer and can vary across Data Feeds. If the TAXII Data Feed content is only to be disseminated to authorized parties, it might make sense to require a subscription. Alternately, Poll Service implementers can allow requests without requiring the Consumer to have established a subscription. This might make sense if the Poll Service supports public feeds to avoid the need to track subscriptions from a large body of users.

Table 8 - TAXII Poll Request Fields

Name	Required?	Multiple?	Description
Feed Name	Yes	No	This field identifies the name of the TAXII Data Feed that is being polled. Each TAXII Data Feed managed by a single Poll Service MUST have a unique Feed Name.
Exclusive Begin Timestamp Label	Yes	No	This field contains a Timestamp Label indicating the beginning of the range of TAXII Data Feed content the requester wishes to receive. This field is exclusive (e.g., the requester is asking for content where the content's Timestamp Label > this field value). A special value is provided by all TAXII Message Binding Specifications to indicate that the requested range has no lower bound.
Inclusive End Timestamp Label	Yes	No	This field contains a Timestamp Label indicating the end of the range of TAXII Data Feed content the requester wishes to receive. This range is inclusive (e.g., the requester is asking for content where the content's Timestamp Label <= this field value). A special value is provided by all TAXII Message Binding Specifications to indicate that the requested range has no upper bound.
Subscription ID	No	No	This field identifies the existing subscription the Consumer wishes to poll. If the Poll Service does not require subscriptions, this field MAY be ignored by the Poll Service. If the Poll Service requires established subscriptions for polling and this field is not present, the Poll Service SHOULD respond with a TAXII Status Message with a status of "Denied".
Content Binding	Yes	Yes	This field indicates the type of content that is requested in the response to this poll. Each Content Binding MUST be a Content Binding ID as defined in the TAXII Content Binding Reference or by a third party. A special value is provided by all TAXII Message Binding Specifications to indicate that the client accepts any type of content.

3.4.9 TAXII Poll Response

This message is sent from a Poll Service in response to a TAXII Poll Request. This message indicates the time bounds within which TAXII Data Feed content was considered in the fulfillment of this request. Note that, as with any content provided by a Producer, the Producer MAY edit or eliminate content for any reason prior to providing it to a Consumer. As such, two Consumers polling the same Poll Service using identical parameters might receive different TAXII Data Feed content. For this reason, the Poll Response Begin Timestamp and End Timestamp fields reflect the range of timestamps the Producer *considers*, but not all content in the considered range is necessarily included in the Poll Response message. Nominally, the timestamp bounds in the Poll Response will be identical to the bounds provided in the Poll Request, with a "No Upper Bound" value replaced by the latest timestamp the Producer considered for inclusion. Under some circumstances, the Producer might provide a different

bound - for example, if the Producer only considered some sub-segment of the Consumer's requested timestamp bounds when producing their response.

Table 9 - TAXII Poll Response Fields

Name	Required?	Multiple?	Description
Feed Name	Yes	No	This field indicates the name of the TAXII Data Feed that was polled. Each TAXII Data Feed managed by a single Poll Service MUST have a unique Feed Name.
Inclusive Begin Timestamp Label	Yes	No	This field contains a Timestamp Label indicating the beginning of the time range this Poll Response covers. A special value is provided by all TAXII Message Binding Specifications to indicate that the Poll Response covers the earliest time for this data feed. This field is inclusive.
Inclusive End Timestamp Label	Yes	No	This field contains a Timestamp Label indicating the end of the time range this Poll Response covers. This field is inclusive.
Subscription ID	No	No	This field contains the Subscription ID for which this content is being provided. This field is only present if this content is being provided as part of an established subscription to a TAXII Data Feed.
Message	No	No	This field contains additional information for the message recipient. There is no expectation that this field be interpretable by a machine; it is instead targeted to human readers.
Content Block	No	Yes	This field contains a piece of content and additional information related to the content. This field MAY appear 0 or more times. See Section 3.5 for the definition of a Content Block.

3.4.10 TAXII Inbox Message

A TAXII Inbox Message is used to push content from one entity to the TAXII Inbox Service of another entity.

Table 10 - TAXII Inbox Message Fields

Name	Required?	Multiple?	Description
Message	No	No	This field contains prose information for the message recipient. This message is not required to be machine readable and is usually a message for a human operator.
Subscription Information	No	No	This field is only present if this message is being sent to provide content in accordance with an existing TAXII Data Feed subscription. It has multiple sub fields:
Feed Name	Yes	No	This field indicates the name of the TAXII Data Feed from which this content is being provided.
Subscription ID	Yes	No	This field contains the Subscription ID for which this content is being provided.

Name	Required?	Multiple?	Description
Inclusive Begin Timestamp Label	Yes	No	This field contains a Timestamp Label indicating the beginning of the time range this Inbox Message covers. A special value is provided by all TAXII Message Binding Specifications to indicate that the Inbox Message covers the earliest time for the subscribed TAXII Data Feed. This field is inclusive.
Inclusive End Timestamp Label	Yes	No	This field contains a Timestamp Label indicating the end of the time range this Inbox Message covers. This field is inclusive.
Content Block	No	Yes	This field contains a piece of content and additional information related to the content. This field MAY appear 0 or more times. See Section 3.5 for the definition of a Content Block.

3.5 TAXII Content Block

A TAXII Content Block contains a piece of content consisting of structured cyber threat information.

Table 11 - TAXII Content Block

Name	Required?	Multiple?	Description
Content Binding	Yes	No	This field contains a Content Binding ID (defined in Section 3.1.6) or nesting expression (defined in Section 5.3) indicating the type of content contained in the Content field of this Content Block. The Content Binding MUST use Content Binding IDs as defined in the TAXII Content Binding Reference or by a third party.
Content	Yes	No	This field contains a piece of content of the type specified by the Content Binding.
Timestamp Label	No	No	This field contains the Timestamp Label associated with this Content Block. It is at the sender's discretion as to whether this is included.
Padding	No	No	This field contains an arbitrary amount of padding for this Content Block. This is typically used to obfuscate the size of the Content Block when the Content is encrypted. This field MUST be ignored when processing a Content Block.
Signature	No	No	This field contains a signature associated with this Content Block. The scope of this field is limited to the Content Block that contains this field.

4 TAXII Message Exchanges

This section describes the TAXII Message Exchanges needed to support the TAXII Services defined earlier. These exchanges only consider TAXII Messages and are agnostic to the network protocols over which those messages travel. In particular, those network protocols might require additional network exchanges prior to transmitting TAXII Messages (e.g., a SSL/TLS handshake) or break a single TAXII

Message into multiple portions that are transmitted independently. The diagrams below represent the conceptual sequence in which TAXII Messages are transmitted and acted upon.

The columns in the exchanges correspond to a TAXII Daemon supporting a specific TAXII Service, as described in the Services section, or a TAXII Client. Note that a single TAXII Daemon might implement multiple TAXII Services. For this discussion we will use a shorthand notation of denoting a TAXII Daemon that supports the ABC Service as an "ABC Daemon". (I.e., a TAXII Daemon that supports the Inbox Service is referred to as an "Inbox Daemon".)

4.1.1 Inbox Exchange

In this exchange, an Inbox Message is transmitted from a TAXII Client to a listening Inbox Daemon. The Inbox Message might be solicited (e.g., a message sent to the recipient as part of a registered subscription) or unsolicited (e.g., an alert sent by some unaffiliated researcher to some public repository). The Inbox Daemon MAY be capable of filtering messages based on the authenticated identity of the sender.

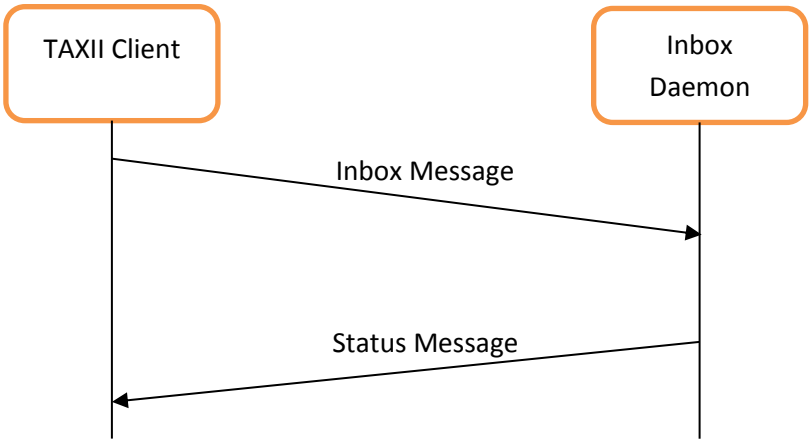

Figure 1 - Inbox Exchange

In this exchange, the TAXII Client sends an Inbox Message to the Inbox Daemon. The Inbox Daemon passes the Inbox Message, along with any authenticated identity information, on to its TAXII Back-end. The TAXII Client receives a Status Message in response from the Inbox Daemon indicating the success or failure of the message exchange. Note that a Status Message of type "Success" indicates only that the Inbox Daemon successfully received and parsed the message. The message might still be discarded by the recipient's TAXII Back-end but the sender receives no indication if this occurs. A Status Message with an error type is used to indicate a problem with the received message.

4.1.2 Discovery Exchange

In this exchange, a TAXII Client requests information about the TAXII Services offered by a particular party. The contacted Discovery Daemon responds with a list of TAXII Services. Note that the Discovery Daemon is not required to reveal all of the TAXII Services of which it is aware to all TAXII Clients.

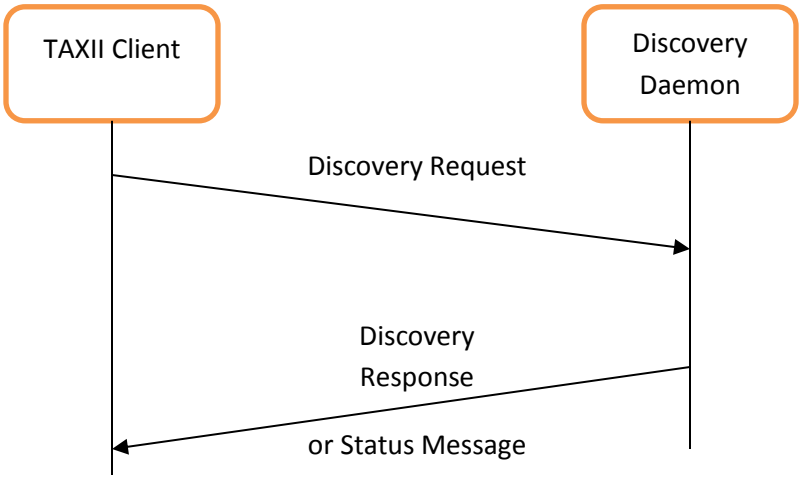

Figure 2 - Discovery Exchange

In this exchange, the TAXII Client sends a Discovery Request to the Discovery Daemon. When the Discovery Daemon receives the Discovery Request Message it can return a TAXII Status Message, indicating an error condition, or pass the relevant information to its TAXII Back-end. Relevant information can include the authenticated identity, if provided. The TAXII Back-end uses this information, along with its own access control policy, to create a list of TAXII Services to be returned or determine that the request will not be fulfilled. (E.g., the request might be denied due to a lack of authorization on the part of the requester.) If the request is honored, the list of TAXII Services is packaged into a Discovery Response which is sent back to the TAXII Client. The TAXII Client receives this message and passes the information to its own TAXII Back-end for processing. The TAXII Status Message MUST only be returned to indicate an error occurred or that the request was denied.

4.1.3 Feed Information Exchange

In this exchange, a TAXII Client requests information about the TAXII Data Feeds available on a Feed Management Daemon. The Feed Management Daemon then responds with a list of available TAXII Data Feeds. The Feed Daemon's response is dictated by its TAXII Back-end, which might consider appropriate access control decisions in composing this response. Note that the Feed Management Daemon is not required to reveal all of the TAXII Data Feeds of which it is aware to all TAXII Clients.

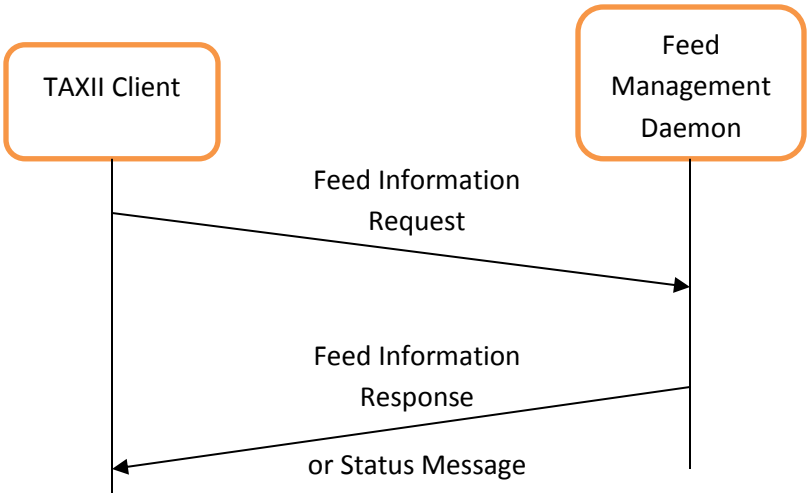

Figure 3 - Feed Information Exchange

In this exchange, the TAXII Client sends a Feed Information Request to the Feed Management Daemon. When the Feed Management Daemon receives the Feed Information Request Message it can return a TAXII Status Message, indicating an error condition, or pass the relevant information to its TAXII Back-end. Relevant information can include the authenticated identity, if any. The TAXII Back-end uses this information, along with its own access control policy, to create a list of feeds to be returned or to determine that the request will not be fulfilled. If the request is honored, the list is packaged into a Feed Information Response that is sent back to the TAXII Client. The TAXII Client receives this message and passes the TAXII Data Feed content to its own TAXII Back-end for processing. The TAXII Status Message MUST only be returned to indicate an error condition or that the request was denied.

4.1.4 Subscription Management Exchange

In this exchange, a client attempts to subscribe, unsubscribe, or request the status of subscriptions to a named TAXII Data Feed by sending a Subscription Management Request to a Feed Management Daemon. The Feed Management Daemon passes the request to its TAXII Back-end, which determines a response. The response is then returned to the TAXII Client.

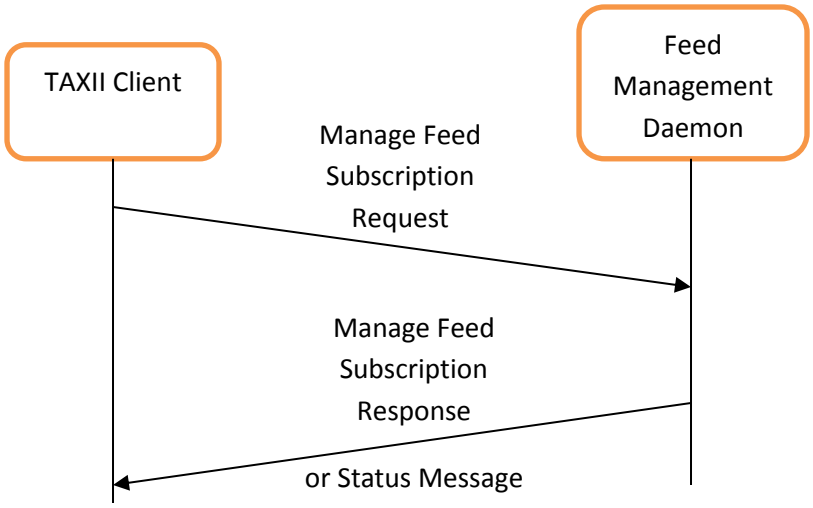

Figure 4 - Subscription Management Exchange

In this exchange, the TAXII Client sends a Manage Feed Subscription Request to the Feed Daemon. The Feed Daemon can immediately return a TAXII Status Message, indicating an error condition, or pass the relevant information to its TAXII Back-end. Relevant information can include the authenticated identity, if any, the parameters for the subscription to be managed, and the action to be taken. The TAXII Back-end uses this information, along with its own access control policy, to determine whether the action is allowed. Depending on this response, the Feed Daemon can return a TAXII Status Message to indicate an error condition or send a Manage Feed Subscription Response. The TAXII Status Message MUST only be returned to indicate an error condition or that the request was denied.

4.1.5 Feed Poll Exchange

This exchange is used by a Consumer to request content from a Producer's TAXII Data Feed. The TAXII Data Feed content is returned to the Consumer in the same exchange. This allows the Consumer to retrieve the TAXII Data Feed content on its own timetable and without needing to field an Inbox Daemon or accept inbound connections. Note that the Poll Daemon is not required to provide all requested content and MAY exclude or alter any content in accordance with its policies.

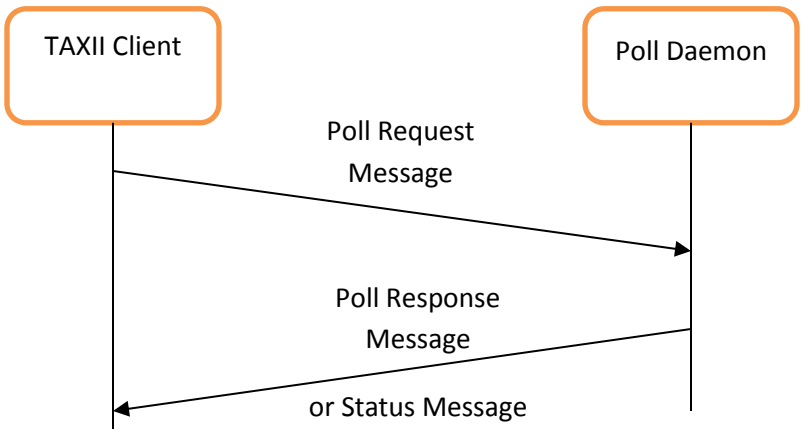

Figure 5 - Feed Poll Exchange

The Consumer's TAXII Client initiates the exchange by sending a Poll Request message to the Producer's Poll Daemon. The Poll Daemon can return an immediate TAXII Status Message, indicating an error condition, or pass the relevant information to its TAXII Back-end. Relevant information can include the Feed Name, Delivery Parameters, Timestamp Labels indicating the interval of information the Consumer is requesting, and the Consumer's authenticated identity, if provided. The TAXII Back-end evaluates this information to determine a response. If the TAXII Back-end decides to honor the request, a Poll Response Message is created encapsulating the provided content. If the TAXII Back-end rejects the request, a TAXII Status message is sent to the client indicating that the request was denied.

In all cases, the TAXII Client receives the appropriate message and passes this information on to its TAXII Back-end for processing. The TAXII Status Message MUST only be returned to indicate an error condition.

5 TAXII Content Handling

This section describes the expected handling of TAXII Content within TAXII Producer Architectures. While the TAXII specifications are agnostic to many aspects of content handling such as how content is stored and access control mechanics, TAXII does impose some requirements on content processing to facilitate compatibility between Producer Architectures.

5.1 Access Control

Many aspects of cyber threat information are considered sensitive by distributing parties. For this reason, some content disseminated using TAXII is likely to be subject to access control protections. TAXII does not stipulate what access controls to impose or how they are implemented, leaving this to individual Producers. However, TAXII does make some assumptions about the overall effect that access control policies can have on content dissemination.

5.1.1 Producers have Full Control over Sharing

Producers have complete discretion as to the information that they share with TAXII Consumers. This includes the ability to redact, alter, or completely hide pieces of TAXII content from TAXII Consumers for any reason. This also includes the ability to hide the presence of TAXII Services in a TAXII Discovery Response and the ability to hide the presence of TAXII Data Feeds in a TAXII Feed Information Response. Moreover, Producers have no obligation to indicate to Consumers that information has been hidden or altered. Even when providing TAXII Status Messages to indicate error conditions, TAXII Producers have discretion as to the amount of detail they provide. In summary, TAXII imposes no requirements on Producers to reveal information if the Producer does not wish to do so.

5.1.2 Changes to Access Levels

If a Consumer's level of access changes with regard to a Data Feed, content that was previously hidden from the Consumer might now be visible. For example, a TAXII Poll Request over a given Timestamp Label range within a TAXII Data Feed might return more (or less) information than was returned by a previous poll request over the same range.

It is outside the scope of TAXII if or how the Consumer's previous requests are updated given their new access rights. TAXII does not include any messages to inform a Consumer that their access rights have changed - informing the client of this is outside the scope of TAXII.

Existing TAXII Subscriptions SHOULD remain valid and active across changes in access level. In other words, if a Consumer has an existing subscription and the Consumer's access rights change, the subscription SHOULD remain operational and the next set of content the Consumer receives uses the Consumers new access rights to determine what content is transmitted.

5.2 Feeds and Content

TAXII Data Feeds are how Producers expose content within a TAXII Architecture. Producers are allowed to assign content to Data Feeds however they wish - feeds can represent communities of users, categories of cyber threat information, or any other grouping the Producer wishes to employ. This section looks at some of the assumptions and requirements surrounding the relationship between TAXII content and TAXII Data Feeds.

5.2.1 TAXII is Content Agnostic

The TAXII specifications do not provide details about the underlying content formats transmitted as content within TAXII. All content formats are a "black-box" as far as TAXII is concerned - none of the behaviors described in TAXII require inspection of any information stored within message content. While TAXII Back-ends can have very different processing paths and requirements for different types of information, TAXII Services, Messages, and Exchanges are agnostic as to the information they convey. This allows TAXII to be usable for a wide array of sharing scenarios.

5.2.2 Content is Static within a Data Feed

As noted earlier in this document, each piece of content is assigned a Timestamp Label when it is added to a TAXII Data Feed. The objective of this is to provide a handle that places that particular piece of

content within an ordering of all of that Data Feed's content. In this way, a Consumer can know that a poll request over a particular range returns all content that will ever appear in that range (modulo the hiding of content for access control reasons). As such (barring changes in the Consumer's access levels) there is no reason to ever re-poll over a given range within a given Data Feed.

For this reason, content MUST NOT be modified after it has been added to a Data Feed. This means that revisions, corrections, and revocations are outside the scope of TAXII. TAXII Messages have no fields to indicate that a new piece of content revises, corrects, or revokes an older piece of content; any such indications need to be expressed within the new content itself, if possible and appropriate.

Producers MAY remove a piece of content from a Data Feed, making it unavailable to further poll requests over a given region. However, deleting content is not an acceptable way to indicate revision, correction, or revocation of that piece of content. Consumers that previously polled over the range that included that piece of content when present have no reason to re-poll over that same range and learn that the content has been removed.

5.3 Content Nesting and Encryption

When conveying TAXII content from a Producer to a Consumer, the Content Binding field in a Content Block indicates the type of content contained in the Content Block's Content field. For example, if the content uses some hypothetical ThreatInfo structure, that ThreatInfo content can be directly ingested by a ThreatInfo -compatible tool once it has been extracted from the Content Block. In other cases, however, content of one type needs to be extracted from content of another type before it can be used. For example, if ThreatInfo content is encrypted, compressed, or otherwise encoded in the Content field, the content of the Content field needs to be processed to extract the ThreatInfo content. TAXII supports multiple methods for indicating the embedding of one form of content inside using the Content Block's Content Binding field.

For the discussion below, suppose a hypothetical "Encryption Structure" exists and is assigned a Content Binding ID of "EncStr". For the ThreatInfo content, assume a ContentBinding ID of "ThreatInfo ". (A real Content Binding ID would include version and format information, but for the sake of generality, the examples below use this simplified ID.)The Encryption Structure contains a field in which one can place a binary blob representing the encrypted form of some content. The following sections describe three ways in which one might use this Encryption Structure to transmit an encrypted content. Note that these examples look at encryption, but other forms of content nesting, such as might be used to support compression, would use identical methods.

5.3.1 Blind Nesting

In Blind Nesting, the Content Binding field identifies only the format of the "outer-most" layer of the Content. In the case of the hypothetical Encryption Structure, this looks something like:

```
Content Binding = EncStr
```

The recipient of a Content Block with this Content Binding knows that they have received an Encryption Structure. However, the Content Binding gives no information as about the content contained within the

Encryption Structure. The recipient needs to determine the nature of the contained content through other means.

5.3.2 Explicit Nesting

In Explicit Nesting, the Content Binding field identifies the type of content at each level of nesting. The Content Binding does this by listing out each Content Binding ID, in order from outer-most to inner-most, separated by a pipe '|' character. In the case of an Encryption Structure containing ThreatInfo content, this might look something like:

```
Content Binding = EncStr|ThreatInfo
```

Explicit nesting makes the type of content the recipient is ultimately receiving clear, although the recipient needs to extract the content from one or more layers of nesting before it can be used. This type of Content Binding value removes any guesswork about the nature of the content within an enclosing structure. On the downside, it also means that an outside observer knows the nature of the content inside the encryption structure, even if they are not able to read that content. This said, explicit nesting is generally viewed as preferable to blind nesting and is recommended over blind nesting when possible.

5.3.3 Content Block Nesting

Instead of containing another content type directly, an outer content type can contain another TAXII Content Block. Each TAXII Message Binding Specification defines its own Content Binding ID to indicate the presence of a Content Block structure within nested content. Assuming a TAXII Message Binding that uses the string "ContentBlock", Content Block nesting looks like the following:

```
Content Binding = EncStr|ContentBlock
```

The following figure demonstrates encryption using Content Block nesting.

A = Payload Block {
    Payload Binding = ThreatInfo
    Payload = ThreatInfo payload
    Signature = Digital signature scoped to **A**
    Padding = ASDFGHJKL...
}

A' = **A**, encrypted and represented in the fictional "Encryption Struct" format

B = Payload Block {
    Payload Binding = EncStr|PayloadBlock
    Payload = A'
    Signature = Digital signature scoped to **B**
}

Figure 6 - Content Block Nesting of ThreatInfo Content

In the above example, **A** represents a Content Block with a piece of ThreatInfo content. The optional Signature field contains a digital signature scoped to this Content Block. The Padding field contains arbitrary data to extend the size of the Content Block.

B represents another Content Block. In this Content Block the content is expressed using the Encryption Struct. For this example, this encrypted material is an encrypted version of the Content Block **A**. The Content Binding field of **B** indicates that the Content field is expressed in the Encryption Struct format and that this structure is wrapping another Content Block. In **B**, the digital signature is scoped to the B Content Block. Note that because **A** is now encrypted, its Padding field obscures the size of the Content field of **A**.

Content Block nesting combines the best aspects of blind and explicit nesting: the type of the inner content is provided explicitly to the recipient once they have extracted and decrypted the Content field from **B** since this information is given explicitly in the Content Binding field of Content Block **A**. At the same time, however, an outside observer can learn nothing about the type of the content being conveyed. In addition, one can see how the Padding field can be used in the inner Content Block to obscure the actual size of the conveyed content.

For the reasons noted above, Content Block nesting is the preferred way of handling content encryption in TAXII over both blind and explicit nesting.

5.3.4 Content Nesting is Disallowed Outside Content Blocks

Content Binding IDs are used in fields outside of a Content Block to indicate content formats that are acceptable within certain contexts, such as within a Data Feed, subscription, or service. Unlike the Content Block's Content Binding field, these fields can contain multiple Content Binding IDs. Nesting expressions (i.e., Content Binding IDs separated by a pipe character) MUST NOT be used in these fields. Instead, when a list of supported content bindings is provided, it indicates that any valid nesting

combination of those bindings is supported. The Content Block is always a supported format and does not need to be listed explicitly.

For example, if a TAXII Poll Request Message indicates the Consumer supports a format W, which is capable of wrapping other content types, as well as formats A and B this indicates support for any valid nesting combination of those formats. E.g., A, B, W, W|A, W|B, W|ContentBlock, W|W|A, W|W|B, etc. are all acceptable formats given the request's supported bindings.

5.4 Sending Requested Content

The ultimate goal of TAXII is to move cyber threat information from a Producer to a Consumer. As noted above, Producers have ultimate control over what gets shared. With that noted, however, Producers do have some obligations to provide the content they are willing to share in certain ways to facilitate Consumer use.

5.4.1 Filtering Content Distribution

Consumers indicate the content bindings they wish to receive, either identifying them when establishing a subscription or when sending a poll request. The list of content bindings indicates the formats the Consumer wishes to receive. Producers SHOULD NOT send content that uses a content binding for which the Consumer did not indicate support. This can mean that certain pieces of content that the Consumer is allowed to receive do not get sent because they can only be expressed using a content binding unsupported by the Consumer.

Note that Producers might not have insight into the nature of content. For example, the content might be encrypted with a key the Producer does not have. (This can happen if the Producer is simply re-sending content send by other parties.) In such a situation, the Producer SHOULD send any content that might be acceptable to a Consumer. For example, consider a Consumer that accepts content wrapped in format W. If the Producer has content expressed in W, but is unaware of what it contains, the Producer SHOULD send the content because the Producer does not know that the Consumer cannot interpret the wrapped content. This means that the Consumer might end up receiving content it is unable to parse. The Consumer MUST NOT treat this as an error condition. (E.g., an Inbox Service that receives content in a format it does not support does not send an "Unsupported Content Binding" Status Message in response.)

5.5 Polling Ranges

When polling, Timestamp Labels are used to specify ranges of content. This section outlines the use of Timestamp Label ranges when formulating responses.

As noted in the section on TAXII Poll Request Messages (Section 3.4.8), poll requests can include an upper and lower Timestamp Label indicating a range of content over which the Consumer is polling. The intent is that this range covers what the Producer "considers" when creating a response. Producers SHOULD honor the Consumer's requested Timestamp Label range when producing a response.

When a Producer "considers" a range, the implication is that all content permitted by access control and similar policies whose Timestamp Label falls between the given bounds is included in the response. In

particular, the Producer MUST include a piece of content in its response if and only if the Producer is willing to share that content with the Consumer and the content's Timestamp Label falls within the Timestamp Label range indicated in the Producer's response message. This means that the Producer's response needs to be complete with regards to the range the Producer indicates in their response.

When determining the bounds of the Timestamp Label range in the Producer's response, the Producer SHOULD use the considered range rather than the actual range of returned content. For example, consider a situation where a Consumer sends a poll request with a lower bound of X. In the polled data feed, there is no content with a Timestamp Label of exactly X, but the content with the next greatest label has a Timestamp Label of Y. The Producer SHOULD use X as the lowest bound in its response because the Producer began its examination of feed content at X, even though it didn't find content until it reached Y. If the Producer were to use Y as its lower bound, the Consumer would not know if there was content with a Timestamp Label between X and Y that it could receive.

Producers MAY send a poll response that indicates a different Timestamp Label range than requested by the Consumer, such as if the user's requested range contained more content than the Producer was willing to send in a single Poll Response.

As noted in Section 3.4.9, Producer MUST always include an upper bound in their poll response message even if the Consumer specified no upper bound. If the Producer's response includes the content with the latest Timestamp Label currently used in the entire feed, the upper bound provided by the Producer MUST be greater than or equal to the Timestamp Label assigned to this last piece of content, and MUST be less than the next Timestamp Label the Producer will assign to the next piece of content added to the Data Feed.

6 Bibliography

[1] The MITRE Corp., "TAXII Overview 1.0," The MITRE Corp., 2013.

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

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

[4] G. Klyne and C. Newman, "RFC 3339 - Date and Time on the Internet: Timestamps," The Internet Engineering Task Force, 2002.

[5] The MITRE Corp., "The TAXII Content Binding Reference," The MITRE Corp., 2013.

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