



TAXII[™] Version 1.1.1. Part 2: Services

Committee Specification Draft 01

06 November 2015

Specification URIs

This version:

<http://docs.oasis-open.org/cti/taxii/v1.1.1/csd01/part2-services/taxii-v1.1.1-csd01-part2-services.docx> (Authoritative)
<http://docs.oasis-open.org/cti/taxii/v1.1.1/csd01/part2-services/taxii-v1.1.1-csd01-part2-services.html>
<http://docs.oasis-open.org/cti/taxii/v1.1.1/csd01/part2-services/taxii-v1.1.1-csd01-part2-services.pdf>

Previous version:

N/A

Latest version:

<http://docs.oasis-open.org/cti/taxii/v1.1.1/taxii-v1.1.1-part2-services.docx> (Authoritative)
<http://docs.oasis-open.org/cti/taxii/v1.1.1/taxii-v1.1.1-part2-services.html>
<http://docs.oasis-open.org/cti/taxii/v1.1.1/taxii-v1.1.1-part2-services.pdf>

Technical Committee:

OASIS Cyber Threat Intelligence (CTI) TC

Chair:

Richard Struse (Richard.Struse@hq.dhs.gov), DHS Office of Cybersecurity and Communications (CS&C)

Editors:

Mark Davidson (mdavidson@mitre.org), MITRE Corporation
Charles Schmidt (cmschmidt@mitre.org), MITRE Corporation
Bret Jordan (bret.jordan@bluecoat.com), Blue Coat Systems, Inc.

Additional artifacts:

This prose specification is one component of a Work Product that also includes:

- *TAXII Version 1.1.1. Part 1: Overview.* <http://docs.oasis-open.org/cti/taxii/v1.1.1/csd01/part1-overview/taxii-v1.1.1-csd01-part1-overview.html>
- *TAXII Version 1.1.1. Part 2: Services* (this document). <http://docs.oasis-open.org/cti/taxii/v1.1.1/csd01/part2-services/taxii-v1.1.1-csd01-part2-services.html>
- *TAXII Version 1.1.1. Part 3: HTTP Protocol Binding.* <http://docs.oasis-open.org/cti/taxii/v1.1.1/csd01/part3-http/taxii-v1.1.1-csd01-part3-http.html>
- *TAXII Version 1.1.1. Part 4: XML Message Binding.* <http://docs.oasis-open.org/cti/taxii/v1.1.1/csd01/part4-xml/taxii-v1.1.1-csd01-part4-xml.html>
- *TAXII Version 1.1.1. Part 5: Default Query.* <http://docs.oasis-open.org/cti/taxii/v1.1.1/csd01/part5-query/taxii-v1.1.1-csd01-part5-query.html>
- XML schemas: <http://docs.oasis-open.org/cti/taxii/v1.1.1/csd01/schemas/>

Related work:

This specification replaces or supersedes:

- *The TAXII Services Specification Version 1.1.*
http://taxiiproject.github.io/releases/1.1/TAXII_Services_Specification.pdf

- *TAXII Content Binding Reference.*
http://taxiiproject.github.io/releases/1.1/TAXII_ContentBinding_Reference_v3.pdf

Declared XML namespaces:

- <http://docs.oasis-open.org/cti/ns/taxii/xml/binding-1.1.1>
- <http://docs.oasis-open.org/cti/ns/taxii/default-query-1.1.1>

Abstract:

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.

Status:

This document was last revised or approved by the OASIS Cyber Threat Intelligence (CTI) TC on the above date. The level of approval is also listed above. Check the "Latest version" location noted above for possible later revisions of this document. Any other numbered Versions and other technical work produced by the Technical Committee (TC) are listed at https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=cti#technical.

TC members should send comments on this specification to the TC's email list. Others should send comments to the TC's public comment list, after subscribing to it by following the instructions at the "[Send A Comment](#)" button on the TC's web page at <https://www.oasis-open.org/committees/cti/>.

For information on whether any patents have been disclosed that may be essential to implementing this specification, and any offers of patent licensing terms, please refer to the Intellectual Property Rights section of the TC's web page (<https://www.oasis-open.org/committees/cti/ipr.php>).

Citation format:

When referencing this specification the following citation format should be used:

[TAXII-v1.1.1-Services]

TAXII[™] Version 1.1.1. Part 2: Services. Edited by Mark Davidson, Charles Schmidt, and Bret Jordan. 06 November 2015. OASIS Committee Specification Draft 01. <http://docs.oasis-open.org/cti/taxii/v1.1.1/csd01/part2-services/taxii-v1.1.1-csd01-part2-services.html>. Latest version: <http://docs.oasis-open.org/cti/taxii/v1.1.1/taxii-v1.1.1-part2-services.html>.

Notices

Copyright © OASIS Open 2015. All Rights Reserved.

All capitalized terms in the following text have the meanings assigned to them in the OASIS Intellectual Property Rights Policy (the "OASIS IPR Policy"). The full [Policy](#) may be found at the OASIS website.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published, and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this section are included on all such copies and derivative works. However, this document itself may not be modified in any way, including by removing the copyright notice or references to OASIS, except as needed for the purpose of developing any document or deliverable produced by an OASIS Technical Committee (in which case the rules applicable to copyrights, as set forth in the OASIS IPR Policy, must be followed) or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by OASIS or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and OASIS DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY OWNERSHIP RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

OASIS requests that any OASIS Party or any other party that believes it has patent claims that would necessarily be infringed by implementations of this OASIS Committee Specification or OASIS Standard, to notify OASIS TC Administrator and provide an indication of its willingness to grant patent licenses to such patent claims in a manner consistent with the IPR Mode of the OASIS Technical Committee that produced this specification.

OASIS invites any party to contact the OASIS TC Administrator if it is aware of a claim of ownership of any patent claims that would necessarily be infringed by implementations of this specification by a patent holder that is not willing to provide a license to such patent claims in a manner consistent with the IPR Mode of the OASIS Technical Committee that produced this specification. OASIS may include such claims on its website, but disclaims any obligation to do so.

OASIS takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on OASIS' procedures with respect to rights in any document or deliverable produced by an OASIS Technical Committee can be found on the OASIS website. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this OASIS Committee Specification or OASIS Standard, can be obtained from the OASIS TC Administrator. OASIS makes no representation that any information or list of intellectual property rights will at any time be complete, or that any claims in such list are, in fact, Essential Claims.

The name "OASIS" is a trademark of [OASIS](#), the owner and developer of this specification, and should be used only to refer to the organization and its official outputs. OASIS welcomes reference to, and implementation and use of, specifications, while reserving the right to enforce its marks against misleading uses. Please see <https://www.oasis-open.org/policies-guidelines/trademark> for above guidance.

Portions copyright © United States Government 2012-2015. All Rights Reserved.

STIX[™], TAXII[™], AND CybOX[™] (STANDARD OR STANDARDS) AND THEIR COMPONENT PARTS ARE PROVIDED "AS IS" WITHOUT ANY WARRANTY OF ANY KIND, EITHER EXPRESSED, IMPLIED, OR STATUTORY, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTY THAT THESE STANDARDS OR ANY OF THEIR COMPONENT PARTS WILL CONFORM TO SPECIFICATIONS, ANY

IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR FREEDOM FROM INFRINGEMENT, ANY WARRANTY THAT THE STANDARDS OR THEIR COMPONENT PARTS WILL BE ERROR FREE, OR ANY WARRANTY THAT THE DOCUMENTATION, IF PROVIDED, WILL CONFORM TO THE STANDARDS OR THEIR COMPONENT PARTS. IN NO EVENT SHALL THE UNITED STATES GOVERNMENT OR ITS CONTRACTORS OR SUBCONTRACTORS BE LIABLE FOR ANY DAMAGES, INCLUDING, BUT NOT LIMITED TO, DIRECT, INDIRECT, SPECIAL OR CONSEQUENTIAL DAMAGES, ARISING OUT OF, RESULTING FROM, OR IN ANY WAY CONNECTED WITH THESE STANDARDS OR THEIR COMPONENT PARTS OR ANY PROVIDED DOCUMENTATION, WHETHER OR NOT BASED UPON WARRANTY, CONTRACT, TORT, OR OTHERWISE, WHETHER OR NOT INJURY WAS SUSTAINED BY PERSONS OR PROPERTY OR OTHERWISE, AND WHETHER OR NOT LOSS WAS SUSTAINED FROM, OR AROSE OUT OF THE RESULTS OF, OR USE OF, THE STANDARDS, THEIR COMPONENT PARTS, AND ANY PROVIDED DOCUMENTATION. THE UNITED STATES GOVERNMENT DISCLAIMS ALL WARRANTIES AND LIABILITIES REGARDING THE STANDARDS OR THEIR COMPONENT PARTS ATTRIBUTABLE TO ANY THIRD PARTY, IF PRESENT IN THE STANDARDS OR THEIR COMPONENT PARTS AND DISTRIBUTES IT OR THEM "AS IS."

Table of Contents

1	Introduction	7
1.1	The TAXII[™] Services Specification	7
1.1.1	TAXII[™] Services Version ID	7
1.2	Terminology	7
1.3	Terms and Definitions	7
1.3.1	TAXII[™] Concepts	7
1.3.2	TAXII[™] Roles	8
1.3.3	TAXII[™] Network Components	8
1.4	Changes from TAXII[™] 1.1	8
1.5	Normative References	8
1.6	Non-Normative References	9
2	TAXII[™] Services	10
2.1	Service Definitions	10
2.1.1	Discovery Service	10
2.1.2	Collection Management Service	10
2.1.3	Inbox Service	11
2.1.4	Poll Service	11
2.2	Service Instances	11
3	TAXII[™] Message Exchanges	12
3.1	Message Overview	12
3.2	Inbox Exchange	12
3.2.1	Pushing Content to TAXII[™] Data Collections	13
3.3	Discovery Exchange	14
3.4	Collection Information Exchange	15
3.5	Subscription Management Exchange	15
3.6	Poll Exchange	16
3.6.1	Multi-Part Poll Exchange	17
3.6.2	Asynchronous Polling	18
4	TAXII[™] Messages	21
4.1	Message Concepts	21
4.1.1	Message IDs	21
4.1.2	Data Collection Names	21
4.1.3	Subscription and Results IDs	21
4.1.4	Timestamp Labels	22
4.1.5	Extended Headers and Status Detail Names	22
4.1.6	Query Format IDs	22
4.1.7	Version IDs, Content Binding IDs, and Content Binding Subtype IDs	23
4.2	TAXII[™] Message Representation Conventions	23
4.3	TAXII[™] Header	24
4.4	TAXII[™] Message Bodies	24
4.4.1	TAXII[™] Status Message	24
4.4.2	TAXII[™] Discovery Request	30
4.4.3	TAXII[™] Discovery Response	30

4.4.4 TAXII[™] Collection Information Request	32
4.4.5 TAXII[™] Collection Information Response.....	32
4.4.6 TAXII[™] Manage Collection Subscription Request.....	35
4.4.7 TAXII[™] Manage Collection Subscription Response.....	38
4.4.8 TAXII[™] Poll Request	40
4.4.9 TAXII[™] Poll Response.....	43
4.4.10 TAXII[™] Inbox Message	45
4.4.11 TAXII[™] Poll Fulfillment Request.....	47
4.5 TAXII[™] Content Block.....	47
5 TAXII[™] Handling.....	48
5.1 Access Control.....	48
5.1.1 Producers have Full Control over Sharing	48
5.1.2 Changes to Access Levels	48
5.2 Data Collections and Content.....	48
5.2.1 TAXII[™] is Content Agnostic.....	48
5.2.2 Data Feeds and Data Sets	49
5.2.3 Directing Inbox Message Content to Data Collections.....	50
5.2.4 Receive-Only Data Collections.....	51
5.3 Content Nesting and Encryption.....	51
5.3.1 Blind Nesting	52
5.3.2 Explicit Nesting	52
5.3.3 Content Block Nesting.....	52
5.3.4 Content Nesting is Disallowed Outside Content Blocks.....	53
5.4 Sending Requested Content.....	53
5.4.1 Targeting Content Requests	53
5.4.2 Paused Subscriptions.....	54
5.5 Query	54
5.5.1 Query Format Specification Requirements	54
5.5.2 General Query Processing	55
6 Conformance	56
Appendix A. Acknowledgements	57
Appendix B. Revision History.....	61

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. Readers are recommended to have familiarity with the relationship between the TAXII specifications, as outlined in the TAXII Overview.

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:oasis:cti:taxii:services:1.1.1

The use of this and other TAXII Version ID strings is described in Section 4.1.7.

1.2 Terminology

The key words “MUST”, “MUST NOT”, “REQUIRED”, “SHALL”, “SHALL NOT”, “SHOULD”, “SHOULD NOT”, “RECOMMENDED”, “MAY”, and “OPTIONAL” in this document are to be interpreted as described in [RFC2119].

1.3 Terms and Definitions

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

1.3.1 TAXII[™] Concepts

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 Collection – A set of structured cyber threat information that can be exchanged using TAXII. Each TAXII Collection has a name that uniquely identifies it among Collections from a given source of Cyber Threat Information. For more on TAXII Data Collection Names, see Section 4.1.2. There are two types of TAXII Data Collections: TAXII Data Feeds and TAXII Data Sets.

TAXII Data Feed – An ordered TAXII Collection. Ordering of records within a TAXII Data Feed is achieved by assigning each piece of content a Timestamp Label. A TAXII Data Feed’s organization allows specific portions of TAXII Data Feeds to be requested (e.g., “Give me all content since I last asked”).

TAXII Data Set – An unordered TAXII Collection.

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 (although portions of TAXII Content may be removed before delivery if the recipient is not allowed to view those portions of the content).

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 4.1.4.

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.

1.4 Changes from TAXII[™] 1.1

From a technical perspective, TAXII 1.1.1 is not substantially different from TAXII 1.1. The intent is that TAXII 1.1.1 represents the smallest set of necessary changes to release TAXII as an OASIS standard. While the technical changes are minimal, the non-technical significance of the TAXII 1.1.1 release is paramount.

The set of technical changes are as follows:

- All references to MITRE and DHS (e.g., The TAXII Services Version for TAXII 1.1 – urn:taxii.mitre.org:services:1.1) have been removed in favor of OASIS conventions.

1.5 Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997. <http://www.ietf.org/rfc/rfc2119.txt>.

- [RFC3986]** T. Berners-Lee, R. Fielding and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", RFC 3986, January 2005. <http://www.ietf.org/rfc/rfc3986.txt>
- [RFC3339]** G. Klyne and C. Newman, "Date and Time on the Internet: Timestamps", RFC 3339, July 2002. <http://www.ietf.org/rfc/rfc3339.txt>

1.6 Non-Normative References

- [XML10]** T. Bray, J. Paoli, C. M. Sperberg-McQueen, E. Maler, and F. Yergeau, "Extensible Markup Language (XML) 1.0 (Fifth Edition)", 2008. <http://www.w3.org/TR/2008/PER-xml-20080205/>

2 TAXIITM Services

TAXII Services represent a set of mechanisms necessary to support one or more TAXII 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.)

2.1 Service Definitions

This section defines the following Services:

- Discovery Service – Provide information about offered TAXII Services.
- Collection Management Service – Support management of TAXII Data Collection subscriptions.
- Inbox Service – Support Producer-initiated pushes of cyber threat information (i.e., push messaging).
- Poll Service - Support Consumer-initiated pulls of cyber threat information (i.e., pull messaging).

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

2.1.1 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 3.3.

2.1.2 Collection Management Service

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

A Subscription can include a query, which restricts delivered content (both push and pull) to only the content in the Collection that meet the criteria specified by the query.

A Collection Management Service implementation MUST support at least one of the Collection Information Exchange or the Subscription Management Exchange, as defined in Sections 3.4 and 3.5, respectively.

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

2.1.3 Inbox Service

The Inbox Service is the mechanism by which a Consumer accepts messages from a Producer in Producer-initiated exchanges. A Consumer can implement this Service in order to receive TAXII Data Collection 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 3.2.

2.1.4 Poll Service

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

Consumers can specify a query when contacting the Poll Service, which restricts the results to only the content in the Collection that meet the criteria specified by the query.

A Poll Service implementation **MUST** support the Poll Exchange, as defined in Section 3.6.

A Poll Service implementation **MAY** support the Multi-Part Poll Exchange, as defined in Section 3.6.1. All TAXII 1.1.1 Clients that communicate with a Poll Service **MUST** support the Multi-Part Poll Exchange.

2.2 Service Instances

This specification makes frequent reference to "service instances". A service instance is defined as a single TAXII Service type over a single protocol binding at a single network address. Note that this definition of a service instances is for bookkeeping rather than a requirement on the actual implementation of a TAXII architecture. For example, in a TAXII architecture it is permissible for a single network address to receive messages for multiple types of TAXII services. However, the TAXII specifications would treat this as involving multiple service instances (one for each supported TAXII Service) despite there being only a single network daemon listening for connections. It is important to remember while TAXII Messages record services using this type-binding-address triple, that the actual implementations of TAXII Services have much greater flexibility.

3 TAXII[™] Message Exchanges

This section describes the TAXII Message Exchanges needed to support the TAXII Services defined above. 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 fragments 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".)

3.1 Message Overview

TAXII Message Exchanges consist only of TAXII Messages as defined in this specification. (See Section 4 for a more detailed description of each TAXII Message.) The TAXII Messages defined in this specification are summarized below.

- TAXII Status Message - Used to indicate an error condition or, in some exchanges, an acknowledgement of message reception.
- TAXII Discovery Request - A request for information about supported TAXII Services.
- TAXII Discovery Response - A response to a TAXII Discovery Request containing information about supported TAXII Services.
- TAXII Collection Information Request - A request for information about supported TAXII Data Collections.
- TAXII Collection Information Response - A response to a TAXII Collection Information Request containing information about supported TAXII Data Collections.
- TAXII Manage Collection Subscription Request - A request to establish a new subscription or manage an existing subscription.
- TAXII Manage Collection Subscription Response - A response to a TAXII Manage Collection Subscription Request indicating the new state of subscriptions to a given TAXII Data Collection.
- TAXII Poll Request - A request for content associated with a TAXII Data Collection.
- TAXII Poll Response - A response to a TAXII Poll Request containing content associated with a TAXII Data Collection.
- TAXII Inbox Message - Used to push content to a recipient.
- TAXII Poll Fulfillment Request - Used to request a delayed result (such as indicated by a "Pending" Status Type) or request another portion of a multi-part result.

3.2 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., a voluntary contribution of content to some recipient). 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. If the Inbox Daemon detects an error that prevents processing of the message (e.g., a malformed message) the Inbox Daemon **MUST** respond with an appropriate Status Message indicating that the exchange failed. Otherwise, the Inbox Daemon passes the Inbox Message, along with any relevant information, on to its TAXII Back-end. The TAXII Daemon **MUST** send a Status Message in response to the Inbox Message, indicating either 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 and that it met TAXII-level requirements (e.g., the content has the right Content Binding ID, is associated with the correct permissions, etc.). The recipient's TAXII Back-end **MAY** still discard the content for any reason and is not required to inform the sender if this happens. A Status Message with a Status Type other than Success is used to indicate a problem with the received message and/or content. If a Status Type other than Success is returned, the Inbox Message recipient **MUST** discard the content received in the Inbox Message. (In other words, if a TAXII Client receives a Status Message with a Status Type other than "Success", they will know that none of the information in the Inbox Message was successfully received.) See Section 4.4.1 for more information on the various Status Types supported by Status Messages and the conditions they indicate.

3.2.1 Pushing Content to TAXII[™] Data Collections

Sometimes the goal of pushing content to a recipient is for the recipient (operating as a TAXII Producer) to add the received content to one or more TAXII Data Collections. For example, this might be the case in a hub-and-spoke architecture where spokes would push content to specific Data Collections on the hub and the hub would then automatically make that content available to other spokes.

There are two ways TAXII can facilitate the automated routing: implicitly and explicitly. In the implicit method, the recipient can simply configure their Inbox Service to automatically add received content for addition to one or more Data Collections. This arrangement can be revealed to senders via descriptions of the Inbox Service itself and via descriptions of the associated Data Collection. The sender simply directs content to the appropriate Inbox Service and needs to take no further action to direct content. Explicit routing is also supported: An Inbox Message in an Inbox Exchange **MAY** identify one or more TAXII Data Collections operated by the message recipient to which the sender wishes content to be added.

It is important to note that TAXII never requires that a recipient of an Inbox Message add any piece of content to a Data Collection. TAXII always permits the recipient to discard the content instead of adding it. Specific sharing agreements established between parties might impose requirements on behavior, but TAXII does not impose such requirements.

The exchange for pushing content to one or more Data Collections is identical to the Inbox Exchange as described above. The sender's TAXII Client sends an Inbox Message to the recipient's Inbox Daemon,

and the Inbox Daemon responds with an appropriate Status Message. However, there are a few additional rules that apply when routing content to Data Collections, be it through the implicit or explicit models:

1. If the recipient's Inbox Daemon returns a Status Message with a Status Type other than "Success", all content in the Inbox Message MUST be discarded.
2. If a Status Message of type Success is returned, all this indicates is that the recipient acknowledges the receipt of the content. The recipient always has discretion to discard content rather than adding it to a Collection, and can do so on a case-by-case basis. If all of the content is discarded, the recipient MAY choose to respond with an error Status Message indicative of the reason for doing this, or it MAY respond with a Success Status Type to indicate that the content was received and processed (the end result of this processing being that the content was discarded). The latter is important since the final decision to add to a Data Collection might only occur after some manual review of the provided content, well after the network exchange has concluded.
3. It is recommended that, for maximum clarity, sharing agreements stipulate that if Destination Collection Names are provided that the content will not be added to Collections other than those named. This gives senders greater control over how content is routed using the explicit model. However, TAXII does not require this behavior to be followed.

For more information on pushing content to named Data Collections see Section 5.2.3.

3.3 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.

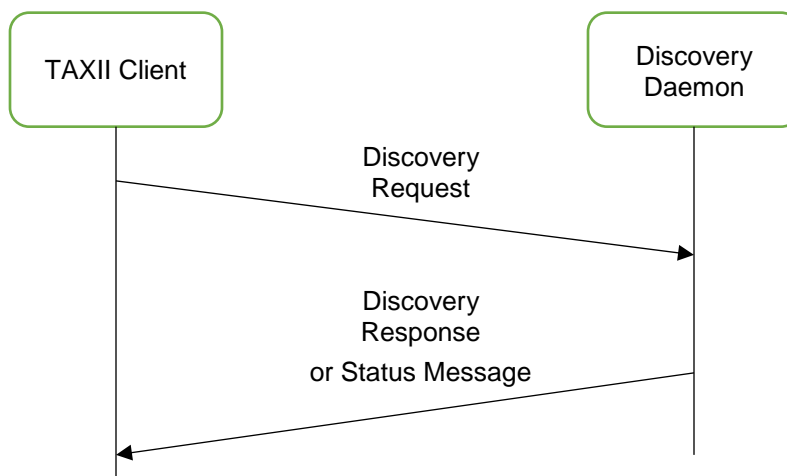


Figure 2 - Discovery Exchange

In this exchange, the TAXII Client sends a Discovery Request to the Discovery Daemon. If the Discovery Daemon detects an error that prevents processing of the message then it MUST respond with an appropriate Status Message indicating that the exchange failed. Otherwise, the Discovery Daemon passes the relevant information to its TAXII Back-end. 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, a list of TAXII Services is packaged into a Discovery Response which is sent back to the TAXII Client. (Not that this list might be 0-length if there are no services the requester is permitted to see.) The TAXII Client receives this message and passes the information to its own TAXII Back-end for processing. If the Discovery Daemon does not respond with a Discovery Response for any reason, the Discovery Daemon MUST respond with a Status Message indicating the

reason that prevented it from returning a successful response. A TAXII Status Message MUST only be returned to indicate an error occurred or that the request was denied.

3.4 Collection Information Exchange

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

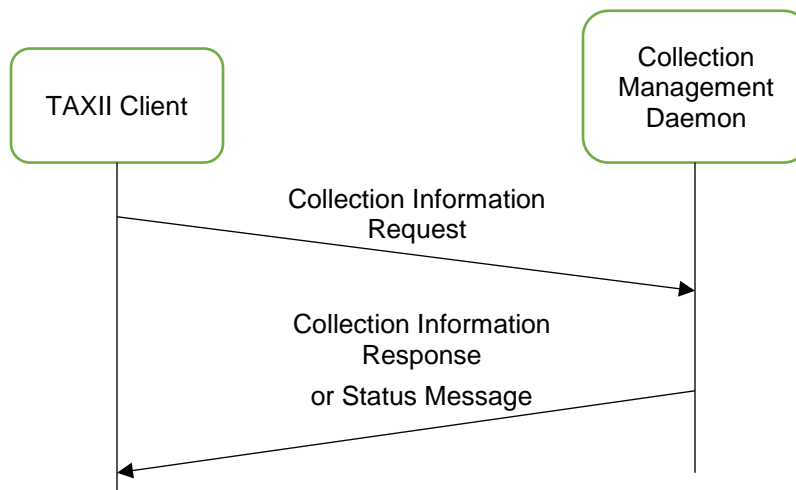


Figure 3 - Feed Information Exchange

In this exchange, the TAXII Client sends a Collection Information Request to the Collection Management Daemon. If the Collection Management Daemon detects an error that prevents processing of the message then it MUST respond with an appropriate Status Message indicating that the exchange failed. Otherwise, the Collection Management Daemon passes the relevant information to its TAXII Back-end. The TAXII Back-end uses this information, along with its own access control policy, to create a list of Data Collections (possibly 0-length) to be returned or to determine that the request will not be fulfilled. If the request is honored, the list is packaged into a Collection Information Response that is sent back to the TAXII Client. The TAXII Client receives this message and passes the TAXII Data Collection content to its own TAXII Back-end for processing. If the request is not fulfilled (i.e., a Collection Information Response is not returned) for any reason, the Collection Management Daemon MUST respond with a Status Message indicating the reason that prevented it from returning a successful response. A TAXII Status Message MUST only be returned to indicate an error condition or that the request was denied.

3.5 Subscription Management Exchange

In this exchange, a client attempts to create a new subscription or change the status of an existing subscription by sending a Manage Collection Subscription Request to a Collection Management Daemon. The Collection Management Daemon passes the request to its TAXII Back-end, which determines a response. The response is then returned to the TAXII Client.

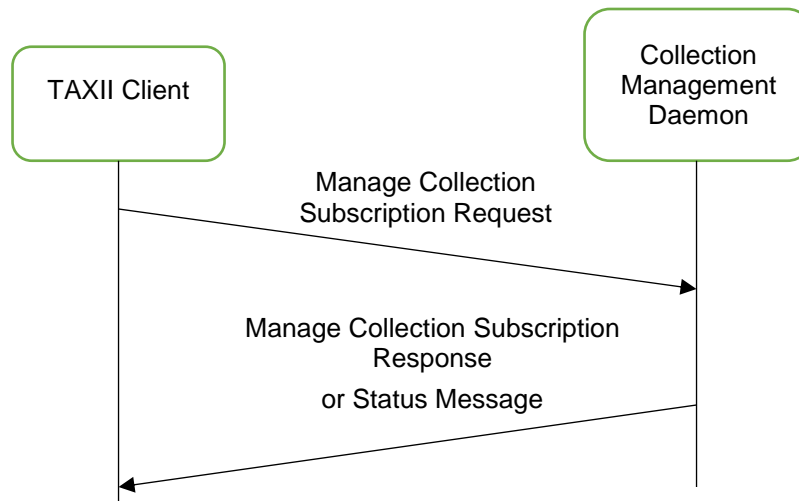


Figure 4 - Subscription Management Exchange

In this exchange, the TAXII Client sends a Manage Collection Subscription Request to the Collection Daemon. If the Collection Management Daemon detects an error that prevents processing of the message then it **MUST** respond with an appropriate Status Message indicating that the exchange failed. Otherwise, the Collection Daemon passes the relevant information to its TAXII Back-end. The TAXII Back-end uses this information, along with its own access control policy, to determine whether the subscription management action is allowed. If the request is allowed, the Collection Management Daemon **MUST** respond with a Manage Collection Subscription Response (even if allowing the request does not change any subscription state). If the request is denied, the Collection Management Daemon **MUST** respond with a Status Message indicating the situation that caused the request to be denied. A TAXII Status Message **MUST** only be returned to indicate an error condition or that the request was denied. The initial request **MUST NOT** result in a change to or the addition of any subscription if a Status Message is returned instead of a Manage Collection Subscription Response.

3.6 Poll Exchange

This exchange is used by a Consumer to request content from a Producer's TAXII Data Collection. The Poll Daemon passes the request to its TAXII Back-end, which determines a response. The response is then returned to the TAXII Client. 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.

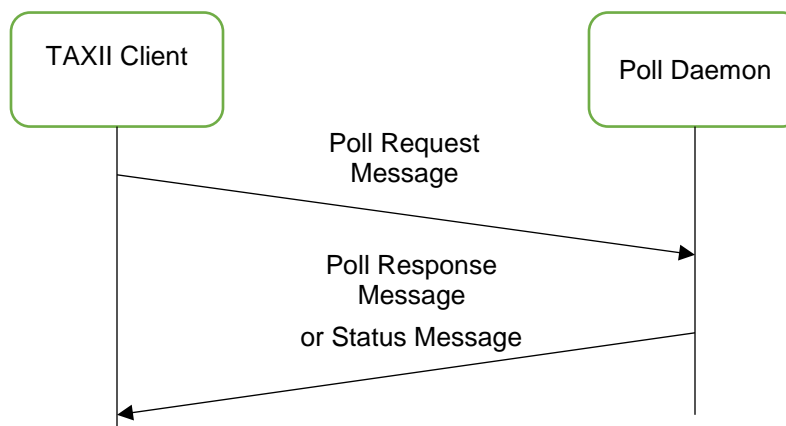


Figure 5 - Poll Exchange

The Consumer's TAXII Client initiates the exchange by sending a Poll Request Message to the Producer's Poll Daemon. If the Poll Daemon detects an error that prevents processing of the message then it MUST respond with an appropriate Status Message indicating that the exchange failed. Otherwise, the Poll Daemon passes the relevant information to its TAXII Back-end. The TAXII Back-end evaluates this information to determine a response. If the TAXII Back-end decides to honor the request and can do so immediately, a Poll Response Message encapsulating the content is created and returned to the TAXII client. Otherwise, the Poll Daemon MUST send a TAXII Status Message to the client indicating that the request was denied or that a result will be provided asynchronously. (See Section 3.6.2 for more on Asynchronous Polling.) In either case, the TAXII Client receives the appropriate message and passes this information on to its TAXII Back-end for processing. The Poll Daemon MUST NOT respond with a TAXII Status Message with a Status Type of "Success".

3.6.1 Multi-Part Poll Exchange

Sometimes the set of content collected in response to a Poll Request can be so large that it is impractical to fit it within a single TAXII Message. A data Producer always has the option of refusing to send such a large result set and instead send an appropriate error using a TAXII Status Message. Sometimes, however, the Producer may still wish to fulfill the Poll Request despite the result size. In this case, the Multi-Part Poll Exchange is employed. The Multi-Part Poll Exchange allows a Consumer to collect the result set of a Poll Response over multiple round-trips.

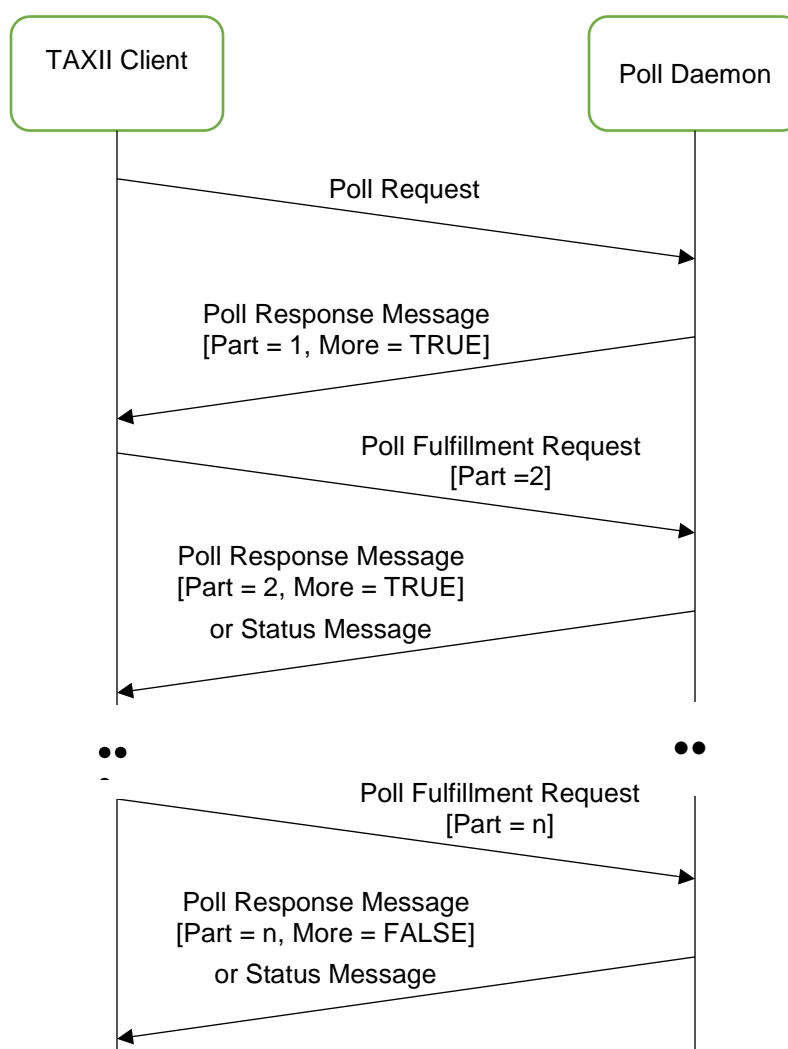


Figure 6 - Multi-Part Poll Exchange

The Multi-Part Poll Exchange begins, as with a normal Poll Exchange, with a Consumer's TAXII Client sending a Poll Request to the Producer's Poll Daemon. (In fact, at the beginning the TAXII Client may be unaware that they will be engaging in a Multi-Part Poll Exchange rather than a Poll Exchange.) In this case, there is no error condition (since otherwise this would just be a simple Poll Exchange with the Poll Daemon returning an immediate Status Message) and the Poll Daemon's TAXII Back-end begins to collect the content to return. At some point the Poll Daemon (or its Back-end) determines that the result set is too large to send in a single TAXII Response. The result set is then divided up so that the result set now consists of multiple parts, such that each part is small enough to be sent in a Poll Response Message. When dividing the result set, records MUST remain whole. (I.e., divisions cannot be in the middle of records but must be made at record boundaries.) Each part is assigned a number starting with 1. A Poll Response encapsulating part 1 of the result set is then created and returned to the TAXII Client. Importantly, the "More" flag in this Poll Response is given a value of TRUE. This flag informs the TAXII Client that there are additional parts of the result set to collect. In addition a result identifier is included in the Poll Response Message so that subsequent parts of the same result set can be identified. This message also contains the Result ID that can be used by the Consumer to identify the result set subsequent messages.

To collect the next part of the result set, the TAXII Client sends a Poll Fulfillment Request including the result identifier and next part number of that result set. If there is a problem with the Poll Fulfillment Request, the Poll Daemon MUST respond with a Status Message indicating the error. Otherwise, the Poll Daemon MUST respond with a Poll Response Message containing the part of the result set identified in the Poll Fulfillment Request. In general, a TAXII Client will request sequential parts of the result set (2, 3, 4, 5, etc.) but is permitted to request any valid part number of the result set. (In most cases, this is done to re-request a part that was corrupted in transit or which was not delivered in an earlier request.) This process repeats until the TAXII Client receives a Poll Response where the More flag is set to FALSE indicating there are no parts of the result set with greater part numbers.

Note that if the TAXII Daemon responds to a particular Poll Fulfillment Request with a Status Message indicating an error, this does not necessarily prevent the TAXII Client from sending another Poll Fulfillment Request and having it succeed. In general, a malformed Poll Fulfillment Request on the part of the TAXII Client should not necessitate that the entire request be started over from the beginning but instead the Producer should allow the TAXII Client to correct its error and continue the collection of result set parts from where it left off. The Producer has discretion as to at what point the parts of the result set are no longer available for collection.

In all individual round-trips, the TAXII Client receives the appropriate response from the Poll Daemon and passes this information on to its TAXII Back-end for processing.

3.6.2 Asynchronous Polling

In addition to immediate response or failure, TAXII supports a feature called Asynchronous Polling. Asynchronous Polling occurs when the recipient of a Poll Request plans to honor the request, but cannot or will not do so immediately. For example, if the compilation of a result set will require hours of time searching non-indexed records, the Producer might choose to utilize Asynchronous Polling to fulfill the request.

Asynchronous Polling begins with the TAXII Client sending a Poll Request as with a normal Poll Exchange. A TAXII Client indicates it is willing to support Asynchronous Polling by setting the Asynch Allowed property of the Poll Request message. (The TAXII Client may be unaware that their Poll Request will require Asynchronous Polling to fulfill but by setting this property the Client is indicating its willingness to participate in such an exchange.) The Poll Daemon receives it and processes it as with the normal Poll Request. At some point it is determined that the result set will not be provided immediately, but the Producer still intends to provide the result set. (If the latter wasn't true, the Producer responds with a Status Message indicating an appropriate error condition and the exchange concludes as a normal Poll Exchange.) If these conditions are met, the Poll Daemon indicates this by responding to the Poll Request with a Status Message with a Status Type of "Pending". This "Pending" Status Message includes an indication as to when the Producer expects that the result will be ready, an identifier by which the specific result set may be indicated, and the manner in which the result set will be made available to the requester. There are two possible ways in which the result set can be provided to the requestor: the

consumer can pull the results or the producer can push the results. The following sections examine each method.

3.6.2.1 Pulling Asynchronous Poll Results

At any time after the TAXII Client receives the "Pending" Status Message it may send a Poll Fulfillment Request Message to the Poll Daemon. The Poll Fulfillment Request includes the identifier provided in the "Pending" Status Message to indicate which result set the request is asking about.

- If there is an error with the Poll Fulfillment Request itself, the Poll Daemon MUST respond with a TAXII Status Message with an appropriate Status Type. The identified result set may still be available if the TAXII Client corrects its error and re-issues a Poll Fulfillment Request.
- If the identified result set is no longer available for collection (e.g., an unexpected error was encountered when generating it, there is no result set with the given result identifier, etc.) then the Poll Daemon MUST respond with a TAXII Status Message with a Status Type of "Asynchronous Poll Error".
- If there is no error but the result set is not yet ready, the Poll Daemon MUST respond with a TAXII Status Message with a Status Type of "Pending". This "Pending" Status Message contains an updated estimate as to when the result will be ready, as well as the other information noted above.
- If the result set is ready and available for collection, the Poll Daemon responds with a Poll Response Message that contains the result set.
 - Note that, if the result set is extremely large, this Poll Response Message may have its "More" field set to true. This would then lead to a Multi-Part Poll Exchange as described in Section 3.6.1, with the difference that the first message from the TAXII Client was a Poll Fulfillment Request rather than a Poll Request.

Once a result set is complete, the Producer has discretion as to how long it is retained. The Producer MAY choose to restrict access to the result set only to the Consumer that issued the request, but could alternatively make it available to any party provided they knew the appropriate result set identifier. (For example, if the specific result set was broadly applicable across a community, the original requester might publish the result set identifier to others and allow them to download the result set immediately, rather than having each party need to make the same, long-duration Poll Request and have the Producer generate the same result set for each party individually.) TAXII can be used in either case, but does not dictate one particular decision.

3.6.2.2 Pushing Asynchronous Poll Results

It is also possible for the result set of an Asynchronous Poll to be pushed to the requester when it is ready. This method involves more architectural complexity on both the part of the Consumer and the Producer. As such, neither Producers nor Consumers are required to support this method, and the method is only used if both Producer and Consumer explicitly agree to do so.

A Consumer indicates its willingness to have results pushed by including information identifying an Inbox Service and how to contact it within its Poll Request Message. This information is ONLY used under an Asynchronous Poll scenario. A Consumer may not know if a particular Poll Request will need to be fulfilled asynchronously at the time the Poll Request is sent. As such, this information is included in the Poll Request as a contingency. (In effect, the Consumer is saying, "If this request is going to be fulfilled asynchronously, I request that result be pushed to the following Inbox Service.")

When the Poll Daemon receives the Poll Request, it only examines these fields identifying an Inbox Service in this request after it has determined that the Poll Request is going to be fulfilled asynchronously. The Producer agrees to push the result set when it is ready to the indicated Inbox Service only if the following requirements are met:

1. The Producer is willing to push this asynchronous result set. (The Producer may refuse to do so for any reason.)
2. The Poll Request contained fields identifying an Inbox Service to receive asynchronous poll results. (I.e., the Consumer has indicated its willingness to receive pushed results.)

3. The Producer is capable of sending an Inbox Message to the indicated Inbox Service using the bindings identified protocol and message bindings.

If all three requirements are met, the Producer indicates signals its intention to push the asynchronous result set when it is ready via a parameter of the "Pending" Status Message it returns to the TAXII Client. The TAXII Client may still send Poll Fulfillment Requests to the Poll Service for the identified result set - usually this is done to update the TAXII Client's understanding of the expected time before the result set is ready. When the result set is complete, the Producer uses a normal Inbox Exchange to deliver this result set to the Consumer, using the identified Inbox Service. If the result set is too large to fit within a single Inbox Message it may be broken into parts, as described in the Multi-Part Poll Exchange in Section 3.6.1 and then delivered via multiple Inbox Exchanges, with each Inbox Message delivering one part of the full result set. (Note that, in this situation, it is not necessary to number the parts.)

The Producer MAY require that the identified Inbox Service authenticate using the same identity as used to issue the original Poll Request. This would prevent malicious parties from issuing Poll Requests with massive results and then directing the result set at someone else's Inbox Service as a way to overwhelm it.

4 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.

4.1 Message Concepts

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

4.1.1 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** conform to URI formatting rules [RFC3986].

4.1.2 Data Collection Names

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

Consumers use Collection Names as handles to a Producer's TAXII Data Collections in their request messages. Note that because Collection 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 Collections with identical Collection Names. For this reason, Consumers need to track both the Collection Name and the associated Producer identity together since the combination of these values is globally unique.

Data Collection names **MUST** conform to URI formatting rules [RFC3986].

4.1.3 Subscription and Results IDs

TAXII Consumers may be able to establish subscriptions to TAXII Data Collections 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 Collection by the same Consumer **MUST NOT** be given the same Subscription ID.

Similarly, in some situations the result set generated from a Poll Request is not returned (or not returned in its entirety) in the Poll Response. (See Multi-Part Poll Exchange and Asynchronous Polling in Sections 3.6.1 and 3.6.2, respectively, for these situations.) When this happens, it is necessary for the Producer to provide an identifier, called a Result ID, that can be used to identify that result set in subsequent exchanges. Not that a Result ID only needs to be created when a Multi-Part Poll Exchange or Asynchronous Polling is employed; in all other cases there is no need to assign a Result ID to a result set. Two result sets for the same TAXII Data Collection **MUST NOT** be given the same Result ID.

Subscription and Result IDs **MUST** conform to URI formatting rules [RFC3986].

4.1.4 Timestamp Labels

Timestamp Labels give a relative ordering to the content within a TAXII Data Feed. Timestamp Labels are used to provide an index into a TAXII Data Feed, making it possible to request content from a TAXII Data Feed that comes “before” and/or “after” a particular Timestamp Label.

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 [RFC3339].
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. While TAXII currently prohibits more than 6 decimal places of fractional second precision within Timestamp Labels, back-end processing SHOULD NOT rely on this since this requirement may be removed in future releases of TAXII.

4.1.5 Extended Headers and Status Detail Names

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 [RFC3986]. 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.

Similarly, TAXII Status Messages can include the Status Detail field, which contains machine-parsable information related to the indicated Status Type. The Status Detail field contains name-value pairs to clearly identify the information conveyed. As with the Extended Headers, names MUST conform to URI formatting rules and SHOULD contain an “authority” part to avoid name collisions.

Values for extended headers and Status Detail values 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.

4.1.6 Query Format IDs

Some TAXII Messages can contain query expressions. Query expressions allow a Consumer to provide a set of criteria that is used to evaluate records within a Data Collection. Records for which the criteria hold true are said to “match” the Query. TAXII provides a Default Query Format, but allows third parties to employ their own Query Format instead. See Section 5.5 for more about the use of Query expressions in TAXII.

All Query Formats expressible in TAXII are assigned a Query Format ID that identifies that Query Format. These Query Format IDs MUST be globally unique. The parties that define the use of a particular Query Format in TAXII are responsible for assigning it a globally unique Query Format ID. All Query Format IDs MUST conform to URI formatting rules [RFC3986]. All such URIs MUST contain an authority component (i.e., a domain name) indicating the authority that controls the meaning of the ID.

4.1.7 Version IDs, Content Binding IDs, and Content Binding Subtype IDs

This document makes references to TAXII "Version IDs", specifically TAXII Services Version IDs, TAXII Protocol Binding Version IDs, TAXII Query Format 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 identifies the format and version of a piece of content in a TAXII Message. Content Binding IDs may appear with Content Binding Subtype IDs. The TAXII Content Binding Reference defines a canonical list of Content Binding IDs and applicable Subtype IDs for a core set of supported content types. Content Binding IDs and Content Binding Subtype ID 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 and Subtype 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 and Subtype IDs. These are used to indicate custom message bindings, custom protocol bindings, or types of content not covered in the TAXII Content Binding Reference. Third parties MUST NOT define alternate TAXII Services Version IDs.

All Version IDs, Content Binding IDs, and Content Binding Subtype IDs MUST conform to URI formatting rules [RFC3986]. All such IDs MUST be globally unique. All ID URIs MUST contain an authority component (i.e., a domain name) indicating the authority that controls the meaning of the ID.

4.2 TAXII[™] Message Representation Conventions

This section contains descriptions 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 [XML10] binding) to account for the intended meaning of a single field as described in this document. 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. Due to changes in capabilities, some fields from TAXII 1.0 were renamed in TAXII 1.1 (and kept for TAXII 1.1.1). In cases where this occurred, the TAXII 1.0 name appears in brackets after the TAXII 1.1.1 name. (E.g., "Collection Name [Feed Name]" indicates that the field name Collection Name was called Feed Name in TAXII 1.0.) When comparing between TAXII versions, treat these names as equivalent.

- **Required?** - Whether the field **MUST** be present in the message. If a field is not required, in most cases its absence indicates that there is no corresponding value or the field is irrelevant in the given situation. However, sometimes absence of an optional field conveys a specific meaning. (I.e., the field has a default value) The descriptions of fields note the cases where this occurs.
- **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. For example, a particular 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.

4.3 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.
Message Body Type	Yes	No	The type of the TAXII Message. Only identifiers for defined TAXII Messages, as defined in Section 4.4, 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 4.1.5.
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.

4.4 TAXII[™] Message Bodies

TAXII Message bodies are used to support specific TAXII Message Exchanges. Each TAXII Message Body Type is described in detail in the following sub-sections.

4.4.1 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 that an error occurred when processing a received TAXII Message. Error

conditions can occur because the request itself was invalid or because the recipient was unwilling or unable to honor the request. The Status Message is also used in the Inbox Exchange (see Section 3.2) to indicate successful reception of an Inbox Message or for Asynchronous Polling (see Section 3.6.2) to indicate a Poll Request will be fulfilled at a later time.

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. Contents of the Status Detail field consist of zero or more name-value pairs. (The details of how these name-value pairs are structured in a particular message binding are provided in the appropriate TAXII Message Binding Specification.) The individual Status Types indicate the standard names and appropriate values for these sub-fields (if any). Values may consist of structured content. Third parties MAY define their own Status Detail sub-fields.
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 Status Type indicates whether any canonical Status Detail name-value pairs are defined for that Status Type. For Status Types for which canonical Status Detail name-value pairs are provided, Status Messages of the indicated Status Type SHOULD provide a Status Detail Field with all of the named subfields. In a few instances the canonical name-value pairs MUST be provided - these cases are noted in the description of the corresponding Status Type. For any Status Type, including Status Types defined by third parties, additional third party name-value pairs may be provided. Each TAXII Message Binding Specification provides structuring details for each of the suggested Status Detail name-value expressions.

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".	
	Status Detail Name	Status Detail Value
	none	

Status Type	Description				
Asynchronous Poll Error	<p>This is used to indicate that a Producer encountered an unexpected error when creating a result set under Asynchronous Polling. (See Section 3.6.2.) As a result, the result set in question is not going to be available to the Consumer.</p> <table> <tr> <th>Status Detail Name</th><th>Status Detail Value</th></tr> <tr> <td colspan="2"><i>none</i></td></tr> </table>	Status Detail Name	Status Detail Value	<i>none</i>	
Status Detail Name	Status Detail Value				
<i>none</i>					
Bad Message	<p>The message sent could not be interpreted by the TAXII Daemon (e.g., it was malformed and could not be parsed).</p> <table> <tr> <th>Status Detail Name</th><th>Status Detail Value</th></tr> <tr> <td colspan="2"><i>none</i></td></tr> </table>	Status Detail Name	Status Detail Value	<i>none</i>	
Status Detail Name	Status Detail Value				
<i>none</i>					
Denied	<p>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 Collection 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".</p> <table> <tr> <th>Status Detail Name</th><th>Status Detail Value</th></tr> <tr> <td colspan="2"><i>none</i></td></tr> </table>	Status Detail Name	Status Detail Value	<i>none</i>	
Status Detail Name	Status Detail Value				
<i>none</i>					
Destination Collection Error	<p>This is used to indicate a problem with the use of the Destination Collection Name field in an Inbox Message. It can indicate either that:</p> <ul style="list-style-type: none"> The recipient of an Inbox Message requires that the sender indicate a Destination Collection Name, but the Inbox Message did not do so. The recipient of an Inbox Message prohibits the sender from dictating a Destination Collection Name, but the Inbox Message had one or more Destination Collection Name fields. <p>See Section 3.6.2 for more on pushing content to Data Collections.</p> <table> <tr> <th>Status Detail Name</th><th>Status Detail Value</th></tr> <tr> <td>Acceptable Destinations</td><td>A list of Data Collection Names to which the sender is permitted to send content. (Specific content may still be rejected from some of these Data Collections for other reasons.) If the specification of Destination Collection Names in an Inbox Message is prohibited, this list is empty.</td></tr> </table>	Status Detail Name	Status Detail Value	Acceptable Destinations	A list of Data Collection Names to which the sender is permitted to send content. (Specific content may still be rejected from some of these Data Collections for other reasons.) If the specification of Destination Collection Names in an Inbox Message is prohibited, this list is empty.
Status Detail Name	Status Detail Value				
Acceptable Destinations	A list of Data Collection Names to which the sender is permitted to send content. (Specific content may still be rejected from some of these Data Collections for other reasons.) If the specification of Destination Collection Names in an Inbox Message is prohibited, this list is empty.				

Status Type	Description				
Failure	<p>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.</p> <table> <tr> <th>Status Detail Name</th><th>Status Detail Value</th></tr> <tr> <td colspan="2"><i>none</i></td></tr> </table>	Status Detail Name	Status Detail Value	<i>none</i>	
Status Detail Name	Status Detail Value				
<i>none</i>					
Invalid Response Part	<p>This Status Type is sent in response to a Poll Fulfillment Request that requests a particular Result Part Number but the result has fewer than that number of parts.</p> <p>The following name-value pair MUST appear in the Status Detail field.</p> <table> <tr> <th>Status Detail Name</th><th>Status Detail Value</th></tr> <tr> <td>Max Part Number</td><td>The largest part number in this multi-part result.</td></tr> </table>	Status Detail Name	Status Detail Value	Max Part Number	The largest part number in this multi-part result.
Status Detail Name	Status Detail Value				
Max Part Number	The largest part number in this multi-part result.				
Network Error	<p>This indicates an error condition at the network level of a TAXII Message exchange. In many cases, a network-level error would occur before the message was passed to a TAXII component, and thus would probably be indicated to the sender using the protocol's native error messages. (E.g., an HTTP error message.) TAXII Message senders need to be able to handle such native protocol errors correctly and should not assume that they will be expressed using this Status Type in a TAXII Status Message. This Status Type is used if there is a need to express this network error in a TAXII-compatible way.</p> <table> <tr> <th>Status Detail Name</th><th>Status Detail Value</th></tr> <tr> <td colspan="2"><i>none</i></td></tr> </table>	Status Detail Name	Status Detail Value	<i>none</i>	
Status Detail Name	Status Detail Value				
<i>none</i>					
Not Found	<p>The request named a target (e.g., a TAXII Data Collection name) that does not exist on the TAXII Daemon.</p> <table> <tr> <th>Status Detail Name</th><th>Status Detail Value</th></tr> <tr> <td>Item</td><td>The target that the TAXII Daemon failed to locate.</td></tr> </table>	Status Detail Name	Status Detail Value	Item	The target that the TAXII Daemon failed to locate.
Status Detail Name	Status Detail Value				
Item	The target that the TAXII Daemon failed to locate.				

Status Type	Description								
Pending	<p>This is sent in response to a Poll Request to indicate that the requested results will be provided at a later time (rather than in a direct Poll Response). It is primarily used in cases where the Poll Request takes more time to process than allowed by the underlying protocol but the Producer still intends to create a result set and make it available.</p> <p>The following name-value pairs MUST appear in the Status Detail field.</p> <table> <tr> <th>Status Detail Name</th><th>Status Detail Value</th></tr> <tr> <td>Estimated Wait</td><td>A positive integer representing the number of seconds expected to be required to produce a result</td></tr> <tr> <td>Result ID</td><td>A value that will be used to identify the result when it is made available</td></tr> <tr> <td>Will Push</td><td>Has a value of TRUE if the Consumer provided Delivery Parameters and the Producer will push results to the indicated Inbox Service when they are ready. Has a value of FALSE otherwise.</td></tr> </table>	Status Detail Name	Status Detail Value	Estimated Wait	A positive integer representing the number of seconds expected to be required to produce a result	Result ID	A value that will be used to identify the result when it is made available	Will Push	Has a value of TRUE if the Consumer provided Delivery Parameters and the Producer will push results to the indicated Inbox Service when they are ready. Has a value of FALSE otherwise.
Status Detail Name	Status Detail Value								
Estimated Wait	A positive integer representing the number of seconds expected to be required to produce a result								
Result ID	A value that will be used to identify the result when it is made available								
Will Push	Has a value of TRUE if the Consumer provided Delivery Parameters and the Producer will push results to the indicated Inbox Service when they are ready. Has a value of FALSE otherwise.								
Polling Not Supported	<p>The requester attempted to create a subscription where the requester only polls for content, but the associated TAXII Data Collection is not available to the requester via polling.</p> <table> <tr> <th>Status Detail Name</th><th>Status Detail Value</th></tr> <tr> <td colspan="2"><i>none</i></td></tr> </table>	Status Detail Name	Status Detail Value	<i>none</i>					
Status Detail Name	Status Detail Value								
<i>none</i>									
Retry	<p>The request cannot be performed at the current time but may be possible in the future. The requested action will not occur until and unless the request is repeated.</p> <table> <tr> <th>Status Detail Name</th><th>Status Detail Value</th></tr> <tr> <td>Estimated Wait</td><td>A positive integer representing the number of seconds expected to be required before a retry of the request might be successful</td></tr> </table>	Status Detail Name	Status Detail Value	Estimated Wait	A positive integer representing the number of seconds expected to be required before a retry of the request might be successful				
Status Detail Name	Status Detail Value								
Estimated Wait	A positive integer representing the number of seconds expected to be required before a retry of the request might be successful								
Unauthorized	<p>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.)</p> <table> <tr> <th>Status Detail Name</th><th>Status Detail Value</th></tr> <tr> <td colspan="2"><i>none</i></td></tr> </table>	Status Detail Name	Status Detail Value	<i>none</i>					
Status Detail Name	Status Detail Value								
<i>none</i>									

Status Type	Description	
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.	
	Status Detail Name	Status Detail Value
	Supported Bindings	A list of acceptable Message Binding IDs.
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.	
	Status Detail Name	Status Detail Value
	Supported Bindings	A list of acceptable Content Binding IDs, including Content Binding Subtype IDs, if applicable.
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.	
	Status Detail Name	Status Detail Value
	Supported Bindings	A list of acceptable Protocol Binding IDs.
Unsupported Query Format	The requester included a Query expression, but the format of the Query Expression was not supported (or the receiving Service does not support Query.)	
	Status Detail Name	Status Detail Value
	Supported Query Formats	A list of acceptable Query Format IDs. If the service does not support Query, this list will be empty.

4.4.1.1 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 non-error conditions.

Status Types defined by a third party MUST conform to URI formatting rules [RFC3986]. In order to avoid accidental name collisions, third party defined Status Types MUST 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.

4.4.2 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.

4.4.3 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. Absence of this field indicates that there are no TAXII Services that can be revealed to the requester.
Service Type	Yes	No	This field identifies the Service Type of this Service Instance (e.g., Poll, Inbox, Collection 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 that can be used to contact 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.
Supported Query	No	Yes	This field indicates that the service supports a particular format of Query expression. This field SHOULD NOT be present for any Service Type other than Collection Management Service or Poll Service; recipients MUST ignore this field for other Service Types. The Query Format subfield identifies the type of query format supported. Other subfields MAY also be present and provide additional support information about the indicated query format - these parameters are identified in the definition of the given query format. (See Section 5.5 for more on Query Format definition.) Multiple instances of this field may appear, but each instance MUST include a different Query Format value. Absence of this field indicates that the identified service does not support the use of Query expressions.

Name		Required?	Multiple?	Description
	Query Format ID	Yes	No	This field contains the Query Format ID that identifies the format of the Supported Query.
	Inbox Service Accepted Content	No	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. 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. Absence of this field when the Service Type field indicates an Inbox Service means that the Inbox Service accepts all content bindings.
	Subtype	No	Yes	This field identifies content binding subtypes of the specified Content Binding. Each Subtype MUST be a Content Binding Subtype ID as defined in the TAXII Content Binding Reference or by a third party. Absence of this field indicates that the Inbox Service accepts all subtypes of the specified Content Binding.
	Available	No	No	This field indicates whether the identity of the requester (authenticated or otherwise) is allowed to access this TAXII Service. This field can indicate that the requester is known to have access, known not to have access, or that access is unknown. Absence of this field indicates that access is unknown.
	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. Recall from Section 2.2 that, in TAXII specifications, a service instance has a single Service Type with a single protocol binding and a single network address for that binding. Each Service Instance field describes a single service instance according to that definition. A Service Instance MAY identify multiple TAXII Message Binding Specifications and (if the TAXII Service is an Inbox Service) multiple content bindings.

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.

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.

4.4.4 TAXII[™] Collection Information Request

This message is sent to a Collection Management Service to request information about the available TAXII Data Collections. The body of this message is empty.

4.4.5 TAXII[™] Collection Information Response

This message is sent in response to a TAXII Collection 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 Collections and can exclude any or all Collections from this response for any reason. For example, the Producer might wish to exclude Collections created for a specific customer from a list of all Collections. As such, different requesters might be given different lists of Collections to their requests to the same Collection Management Service.

Table 5 - TAXII[™] Collection Information Response Fields

Name		Required?	Multiple?	Description
Collection Information <i>[Feed Information]</i>		No	Yes	This field MAY appear any number of times (including 0), each time identifying a different TAXII Data Collection. It has several sub-fields.
	Collection Name <i>[Feed Name]</i>	Yes	No	This field contains the name by which this TAXII Data Collection is identified.
	Collection Type	No	No	This field indicates whether this Data Collection is a Data Feed (ordered Collection) or a Data Set (unordered Collection). Absence of this field denotes that this Collection is a Data Feed.
	Collection Description <i>[Feed Description]</i>	Yes	No	This field contains a prose description of this TAXII Data Collection. This field might also explain how to gain access to this TAXII Data Collection if out-of-band actions are required. (E.g., requires purchase of a contract, requires manual approval, etc.)
	Collection Volume	No	No	This field indicates the typical number of records added to this Data Collection daily. This represents a "typical" value and the producer is under no obligation to keep the Data Collection volume at the given level.
	Supported Content	No	Yes	This field contains Content Binding IDs indicating which types of content might be found in this TAXII Data Collection. Each Supported Content value MUST be a Content Binding ID as defined in the TAXII Content Binding Reference or by a third party. Absence of this field indicates that this Data Collection supports all types of content.
	Subtype	No	Yes	This field identifies content binding subtypes of the specified Supported Content binding. Each Subtype MUST be a Content Binding Subtype ID as defined in the TAXII Content Binding Reference or by a third party. Absence of this field indicates that this Data Collection supports all subtypes of the specified Supported Content binding.

Name		Required?	Multiple?	Description
	Available	No	No	This field indicates whether the identity of the requester (authenticated or otherwise) is allowed to access this Collection. (Access could imply the ability to subscribe and/or the ability to send Poll Requests.) This field can indicate that the requester is known to have access, known not to have access, or that access is unknown. Absence of this field indicates that access is unknown.
	Push Method	No	Yes	This field identifies the protocols that can be used to push content from this Data Collection via a subscription and/or for pushed results of Asynchronous Polling. This field MAY appear multiple times if content from this TAXII Data Collection can be pushed via multiple protocols. This field has multiple sub-fields. Absence of this field indicates that content from this Data Collection cannot be pushed to a Consumer using TAXII.
	Push Protocol	Yes	No	This field identifies a protocol binding that can be used by the Producer to push content from this Data Collection to a Consumer's 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 by the Producer to push content from this Data Collection 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.
	Polling Service Instance	No	Yes	This field identifies the bindings and address a Consumer can use to interact with a Poll Service instance that supports this TAXII Data Collection. This field MAY appear multiple times if multiple Poll Services support this TAXII Data Collection. This field has multiple sub-fields. Absence of this field indicates that this Data Collection cannot be polled using TAXII.
	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 that can be used to contact the TAXII Daemon hosting this Poll Service instance. This field MUST use a format appropriate to the Poll Protocol field value.

Name		Required?	Multiple?	Description
	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 that can be used to contact the TAXII Daemon hosting the Collection Management Service that can process subscription requests for this TAXII Data Collection. Absence of this field indicates that there is not a TAXII Service that processes subscription requests for this Collection. In that case subscriptions, if supported, would need to be established by mechanisms other than TAXII. In the case of alternative subscription methods, the Collection Description field could provide procedures for subscribing.
	Subscription Protocol	Yes	No	This field identifies the protocol binding supported by this Collection 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 that can be used to contact the TAXII Daemon hosting this Collection 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 Collection 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.
	Receiving Inbox Service	No	Yes	This field identifies the bindings and address of an Inbox Service to which content can be pushed to have it added to the given Data Collection. This field MAY appear multiple times if multiple Inbox Services may receive content for this TAXII Data Collection. If this field is absent, the Consumer cannot use TAXII Messages to request that content to be added specifically to this Data Collection. Note that content sent to this Inbox Service MAY still be rejected by the recipient for any reason instead of adding it to the indicated Data Collection.
	Inbox Protocol	Yes	No	This field identifies the protocol binding supported by this 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.

Name		Required?	Multiple?	Description
	Inbox Address	Yes	No	This field identifies the address that can be used to contact the TAXII Daemon hosting this Inbox Service instance. This field MUST use a format appropriate to the Inbox Protocol field value.
	Inbox Message Binding	Yes	Yes	This field identifies the message bindings supported by this Inbox 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.
	Supported Content	No	Yes	This field contains Content Binding IDs indicating that the indicated Inbox Service only accepts content using specific content bindings. Each Supported Content value MUST be a Content Binding ID as defined in the TAXII Content Binding Reference or by a third party. Absence of this field indicates that the Inbox Service supports all content bindings supported by the Data Collection.
	Subtype	No	Yes	This field identifies content binding subtypes of the specified Supported Content binding. Each Subtype MUST be a Content Binding Subtype ID as defined in the TAXII Content Binding Reference or by a third party. Absence of this field indicates that the Inbox Service supports all subtypes of the given Supported Content binding supported by the Data Collection.

4.4.6 TAXII[™] Manage Collection Subscription Request

This message is used to establish a new subscription or manage an existing subscription. The Collection Management Service responds with a TAXII Manage Collection 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 Collection Subscription Request Fields

Name	Required?	Multiple?	Description
Collection Name [Feed Name]	Yes	No	This field identifies the name of the TAXII Data Collection to which the action applies.

Name		Required?	Multiple?	Description
Action		Yes	No	<p>This field identifies the requested action to take. The action MUST be one of the following:</p> <ul style="list-style-type: none"> ○ SUBSCRIBE - Request a subscription to the named TAXII Data Collection. ○ UNSUBSCRIBE - Request cancellation of an existing subscription to the named TAXII Data Collection. ○ PAUSE – Suspend delivery of content for the identified subscription. ○ RESUME – Resume delivery of content for the identified subscription. ○ STATUS - Request information on subscriptions the requester has established for the named TAXII Data Collection. No subscription state is changed in response to this action.
Subscription ID		Per Action	No	<p>This field contains the ID of a previously created subscription. For messages where the Action field is UNSUBSCRIBE, PAUSE, or RESUME, this field MUST be present. For messages where the Action field is SUBSCRIBE, this field MUST be ignored. For messages where the Action field is STATUS, this field MAY be present.</p>
Subscription Parameters		Yes, if and only if the value of the Action field is SUBSCRIBE	No	<p>This field contains multiple subfields that indicate various aspects of the requested subscription. This field MUST be included if and only if the Action of this message is SUBSCRIBE and MUST be ignored for all other Action values.</p>
	Response Type	No	No	<p>This field identifies the response type that is being requested as part of this subscription. The Response Type MUST be one of the following:</p> <ul style="list-style-type: none"> • FULL – Messages sent in fulfillment of this request are requested to contain full content. • COUNT ONLY – The requester is requesting that messages sent in fulfillment of this subscription only contain count information (i.e., content is not included). <p>Absence of this field indicates a request for FULL responses.</p>
	Content Binding	No	Yes	<p>This field contains Content Binding IDs indicating which types of contents the Consumer requests to receive for this subscription. Multiple Content Binding IDs may be specified. This field MUST contain 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 Collection, a Status Message with a status of 'Unsupported Content Binding' SHOULD be returned. Absence of this field indicates that all content bindings are accepted.</p>

Name		Required?	Multiple?	Description
	Subtype	No	Yes	This field identifies content binding subtypes of the specified Content Binding. Each Subtype MUST be a Content Binding Subtype ID as defined in the TAXII Content Binding Reference or by a third party. Absence indicates that all subtypes of the specified Content Binding are accepted.
	Query	No	No	This field contains a query expression associated with this subscription request. If the subscription request is successful, only content that matches the query expression should be sent in fulfillment of the subscription. The query expression may be structured; the specific structure used for the query expression is identified in the Query Format field.
	Query Format	Yes	No	This field contains a Query Format ID that identifies the format of the query expression that appears within the Query field.
Delivery Parameters		No	No	This field identifies the parameters used to push content to the Consumer in fulfillment of a subscription. This field is only meaningful if the Action field is equal to SUBSCRIBE and is ignored for all other Action values. Absence of this field for a SUBSCRIBE action indicates that the requester is not requesting pushed content and will instead poll for subscription content use a Poll Service. In this case, if the TAXII Data Collection cannot be polled, a Status Message with a status of 'Polling Not Supported' SHOULD be returned.
	Inbox Protocol	Yes	No	This field identifies the protocol to be used when pushing TAXII Data Collection content to a Consumer's TAXII Inbox Service implementation. If the Data Collection 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.
	Inbox Address	Yes	No	This field identifies the address that can be used to contact the TAXII Daemon hosting the Inbox Service to which the Consumer requests content for this TAXII Data Collection 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	No	This field identifies the message binding to be used to send pushed content for this subscription. If the TAXII Data Collection does not support the Delivery Message Binding, a Status Message with a status of 'Unsupported Message Binding' SHOULD be returned. The Delivery Message Binding MUST be a TAXII Message Binding Version ID as defined in a TAXII Message Binding Specification or by a third party.

Manage Collection Subscription 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 Collections where the requested Collection Name does not correspond to an existing Collection 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 Collection by the identified Consumer SHOULD be treated as a successful attempt to unsubscribe and result in a TAXII Manage Collection 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, Query Format, or Content Binding (and Subtype(s), if present) of the subscription to be created is not supported SHOULD result in an appropriate TAXII Status Message (normally "Unsupported Protocol Binding", "Unsupported Message Binding", "Unsupported Query Format", 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 Collection Name, Subscription Parameters, and Delivery Parameters) SHOULD result in a TAXII Manage Collection Subscription Response that returns that existing subscription's Subscription ID without changing existing subscriptions. That is, the Collection Management Service SHOULD NOT create exact duplicates of existing subscriptions, but the client SHOULD be informed that the requested subscription is established.
6. Attempts to pause (PAUSE action) an already paused subscription SHOULD result in a TAXII Manage Collection Subscription Response (indicating success) without changing the existing subscription status. Likewise, attempts to resume (RESUME action) delivery of a subscription that was not in a paused state should result in a TAXII Manage Collection Subscription Response without changing the existing subscription status. In other words, attempts to set the pause-resume state to the current state should appear to be successful.
7. Attempts to pause (PAUSE action), resume (RESUME action), or collect status about (STATUS action) a non-existent subscription (i.e., a Subscription ID value that does not correspond to an existing subscription for the named Data Collection for the requesting party) SHOULD result in a TAXII Status Message of type "Not Found".

4.4.6.1 STATUS Actions

The purpose of the STATUS action is to allow a subscriber to retrieve information about their existing subscriptions to a particular Data Collection. This might be necessary if the subscriber's records are lost and it is not clear which subscriptions have been created, what state they are in, and/or how those subscriptions are configured.

STATUS Actions do not alter subscriptions on the Collection Management Service. If a Manage Collection Subscription Request with a STATUS Action does not include a Subscription ID value, this represents a request for information about all established subscriptions (paused or active) managed by this Collection Management Service. If a Subscription ID value is provided, this represents a request for information only about the named subscription. The Manage Collection Subscription Response Message sent in response to either type of STATUS request contain descriptions of each applicable subscription and its state (usually in the form of copies of the relevant fields from the Manage Collection Subscription Request that contained the SUBSCRIBE Action that established this subscription).

4.4.7 TAXII[™] Manage Collection Subscription Response

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

Table 7 - TAXII[™] Manage Collection Subscription Response Fields

Name		Required?	Multiple?	Description
Collection Name <i>[Feed Name]</i>		Yes	No	This field identifies the name of the TAXII Data Collection to which the action applies.
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		Per Action in the Manage Collection Subscription Request	Yes	This field contains information about existing subscriptions by the requester to the given TAXII Data Collection. 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.
	Status	No	No	<p>This field contains the status of the Subscription. Possible status values are:</p> <ul style="list-style-type: none"> • Active - The subscription is established and active • Paused - The subscription is established but currently in a paused state • Unsubscribed - The subscription has been removed (would only appear in response to an UNSUBSCRIBE Action) <p>If this field is absent, treat it as having a value of Active.</p>
	Subscription Parameters	Per Action in the Manage Collection Subscription Request	No	This field contains a copy of the Subscription Parameters of the Manage Collection Subscription Request message that established this subscription. This field MUST be present if this message is in response to a request with an Action field value of STATUS. This field MAY be present when responding to any other Action type.
	Response Type	No	No	These fields all contain copies of the corresponding fields in the Manage Collection Subscription Request Message that established this subscription. A given field will only be present here if it was present in that Request Message.
	Content Binding	No	Yes	
	Subtype	No	No	
	Query	No	No	
	Query Format	Yes	No	

Name		Required?	Multiple?	Description
	Delivery Parameters	No	No	This field contains a copy of the Delivery Parameters (if present) of the Manage Collection Subscription Request Message that established this subscription. This field is present if and only if the Producer is willing and able to push content to the indicated Inbox Service in fulfillment of the established subscription. (It does not matter whether the subscription is currently in a PAUSED state.)
	Inbox Protocol	Yes	No	These fields all contain copies of the corresponding fields in the Manage Collection Subscription Request Message that established this subscription.
	Inbox Address	Yes	No	
	Delivery Message Binding	Yes	No	
	Poll Instance	No	Yes	Each Poll Instance represents an instance of a Poll Service that can be contacted to retrieve content associated with the named subscription. Its subfields indicate where Poll Request Messages can be sent for the given subscription. Multiple instances of this field may be present if there are multiple Poll Services that can be contacted for content for this subscription. If this field is absent, this indicates that polling for subscription content is not supported via TAXII.
	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 that can be used to contact 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.

4.4.8 TAXII[™] Poll Request

This message is sent from a Consumer to a TAXII Poll Service to request that data from the TAXII Data Collection be returned to the Consumer. Poll Requests are always made against a specific TAXII Data Collection. Whether or not the Consumer needs an established subscription to that TAXII Data Collection in order to receive content is left to the Producer and can vary across Data Collections.

Table 8 - TAXII[™] Poll Request Fields

Name	Required?	Multiple?	Description
------	-----------	-----------	-------------

Name		Required?	Multiple?	Description
Collection Name [Feed Name]		Yes	No	This field identifies the name of the TAXII Data Collection that is being polled.
Exclusive Begin Timestamp Label		No	No	This field contains a Timestamp Label indicating the beginning of the range of TAXII Data Feed (i.e., ordered TAXII Data Collection) content the requester wishes to receive. The receiving TAXII Poll Service MUST ignore this field if the named TAXII Data Collection is a Data Set (i.e., an unordered TAXII Data Collection). This field is exclusive (e.g., the requester is asking for content where the content's Timestamp Label is greater than this field value). Absence of this field when polling a Data Feed indicates that the requested range has no lower bound.
Inclusive End Timestamp Label		No	No	This field contains a Timestamp Label indicating the end of the range of TAXII Data Feed content the requester wishes to receive. The receiving TAXII Poll Service MUST ignore this field if the named TAXII Data Collection is a Data Set. This range is inclusive (e.g., the requester is asking for content where the content's Timestamp Label is less than or equal to this field value). Absence of this field when polling a Data Feed indicates that the requested range has no upper bound.
Subscription ID		Exactly one of Subscription ID or Poll Parameters MUST be present	No	This field identifies the existing subscription the Consumer wishes to poll. 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".
Poll Parameters			No	This field contains multiple subfields that indicate the content to return in the Poll Response. This field MUST NOT be present if a Subscription ID is provided; if a Subscription ID is provided, the corresponding information from the subscription is used instead.
	Response Type	No	No	<p>This field identifies the response type that is being requested. The Response Type MUST be one of the following:</p> <ul style="list-style-type: none">FULL – Messages sent in fulfillment of this request are requested to contain full content.COUNT ONLY – The requester is requesting that messages sent in fulfillment of this subscription only contain count information (i.e., content is not included). <p>Absence of this field indicates a request for FULL responses.</p>

Name		Required?	Multiple?	Description
	Content Binding	No	Yes	This field contains Content Binding IDs indicating which types of contents the Consumer requests to receive. Multiple Content Binding IDs may be specified. This field MUST contain 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 Collection, a Status Message with a status of 'Unsupported Content Binding' SHOULD be returned. Absence of this field indicates that all content bindings are accepted.
	Subtype	No	No	This field identifies content binding subtypes of the specified Content Binding. Each Subtype MUST be a Content Binding Subtype ID as defined in the TAXII Content Binding Reference or by a third party. Absence of this field indicates that all subtypes of the specified Content Binding are accepted.
	Query	No	No	This field contains a query expression. Only content that matches the query expression should be sent in response to this message. The query expression may be structured; the specific structure used for the query expression is identified in the Query Format field.
	Query Format	Yes	No	This field contains a Query Format ID that identifies the format of the query expression that appears within the Query field.
	Allow Asynch	No	No	This field indicates whether the Consumer is willing to support Asynchronous Polling. If this value is FALSE, the response MUST NOT respond with a Status Message with Status Type of "Pending". Absence of this field should be treated as indicating a value of FALSE. For more information on Asynchronous Polling, see Section 3.6.2.
	Delivery Parameters	No	No	This field identifies how to push Asynchronous Poll Results to an Inbox Service specified by the poll requestor if the requestor wishes this to happen. This field MUST NOT be present if Allow Pending is absent or has a value of FALSE. If this field is absent but Allow Pending has a value of TRUE, this indicates that the Consumer will pull any Asynchronous Poll results rather than having them pushed. The Poll Service ignores this field if it is able to include results in a Poll Response Message. (Unsupported sub-field values should not lead to error Status Messages if the Delivery Parameters are ignored.) The Poll Service also ignores this field if it is not willing to push Asynchronous Poll Results to a Consumer.
	Inbox Protocol	Yes	No	This field identifies the protocol to be used when pushing Asynchronous Poll Results to a Consumer's TAXII Inbox Service implementation. 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	No	This field identifies the address that can be used to contact the TAXII Daemon hosting the Inbox Service to which Asynchronous Poll Results may be delivered. The address MUST be of the appropriate type for the network protocol identified in the Inbox Protocol field.
	Delivery Message Binding	Yes	No	This field identifies the message binding to be used to send pushed Asynchronous Poll Results. The Delivery Message Binding MUST be a TAXII Message Binding Version ID as defined in a TAXII Message Binding Specification or by a third party.

The Delivery Parameters field and subfields MAY be included in a Poll Request in support of using pushed messages to fulfill Asynchronous Polling and is discussed in Section 3.6.2.2. In preparation for this possibility, a Poll Request Message MAY contain Delivery Parameters that indicate how Asynchronous Poll Results may be pushed to an Inbox Service designated by the requestor once those results are ready. The Delivery Parameters are only applicable if:

1. Asynchronous Polling is supported by the Poll Service
2. Asynchronous Polling is necessary due to the long processing time needed to service the Poll Request
3. The Poll Service is willing to push Asynchronous Poll Results

The Poll Service should only act on the values of these fields if all of those conditions are met and should ignore them otherwise. In particular, even unsupported values in these fields should not result in Status Messages with error Status Types unless all three of the noted preconditions are met.

4.4.9 TAXII[™] Poll Response

This message is sent from a Poll Service in response to a TAXII Poll 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 Collection content.

If the named TAXII Data Collection is a TAXII Data Feed, the message indicates the time bounds within which TAXII Data Feed content was considered in the fulfillment of this request. As noted, content may be hidden from some Consumers, so 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
Collection Name [Feed Name]	Yes	No	This field indicates the name of the TAXII Data Collection that was polled.
Subscription ID	No	No	This field contains the Subscription ID for which this content is being provided. Absence of this field indicates that this content is not being provided as part of an established subscription to a TAXII Data Collection.

Name		Required?	Multiple?	Description
Exclusive Begin Timestamp Label	No; At most one of these fields may appear	No	No	These fields serve the same purpose. Use of the Inclusive Begin Timestamp Label field is deprecated but retained for backwards compatibility with TAXII 1.0. Both fields MUST NOT appear together in the same message. Either field contains a Timestamp Label indicating the beginning of the time range this Poll Response covers. (One field provides an exclusive value; the other provides an inclusive value.) Absence of either field indicates that the Poll Response covers the earliest time for this TAXII Data Feed. The fields MUST NOT be included if the named TAXII Data Collection is a Data Set.
Inclusive Begin Timestamp Label			No	
Inclusive End Timestamp Label	Required if for a Feed; prohibited otherwise	No	No	This field contains a Timestamp Label indicating the end of the time range this Poll Response covers. This field is inclusive. This field MUST be present if the named Data Collection is a Data Feed. This field MUST NOT be present if the named Data Collection is a Data Set.
More	No	No	No	This field contains a boolean value. If the field value is TRUE, this indicates there are additional parts remaining of a larger result set. If the field value is FALSE, this indicates that there are no parts of the result set with higher Result Part Numbers. If this field is absent, treat that as equivalent to a value of FALSE.
Result ID	No	No	No	This field contains a Result ID that can be used in Poll Fulfillment Requests to identify other parts of this result set. This field MUST be present if the More field is set to TRUE.
Result Part Number	No	No	No	This field contains an integer indicating the part of the result set contained in this Poll Response Message. Each part of a multi-part response is assigned a sequential integer starting with 1. (As such, the response to the initial Poll Request would have a 1 for this field.) If this field is absent, treat the field as having a value of 1.
Record Count	No	No	No	Indicates the number of applicable records for the given Poll Request, which MUST be greater than or equal to the number of content records returned in this message's Content Block(s). This field SHOULD be present in all Poll Response messages.
Partial Count	No	No	No	This field indicates whether the provided Record Count is the exact number of applicable records, or if the provided number is a lower bound and there may be more records than stated. The field contains a boolean value. A value of TRUE indicates that the actual number of matching records may be greater than the value that appears in the Record Count field. A value of FALSE indicates that the Record Count is an exact count of applicable records. If this field is absent, treat the field as having a value of FALSE.

Name	Required?	Multiple?	Description
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 4.5 for the definition of a Content Block.

Note that TAXII 1.1.1 includes two fields to indicate the beginning Timestamp Label value. These fields are only applicable when providing content from a Data Feed; it MUST NOT be the case that either field is present when providing content from a Data Set. Absence of both fields when providing content from a Data Feed indicates that the Poll Response covers the earliest records for the specified Data Feed. When providing a lower Timestamp Label bound for a TAXII Data Feed the TAXII Poll Response Message MUST use exactly one of these fields to indicate the lower bound of the content considered for the response. Use of the Inclusive Begin Timestamp Label field is deprecated, but retained for backwards compatibility. All TAXII 1.1.1 Poll Service implementations SHOULD be able to operate correctly if either field is present. All TAXII 1.1.1 implementations that send TAXII Poll Response Messages SHOULD use the Exclusive Begin Timestamp Label field in their response unless the requestor indicated that it does not accept responses using this field (as indicated by the supported Message Binding version(s)).

4.4.10 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
Destination Collection Name	No	Yes	This field indicates the name of the TAXII Data Collection(s) to which this message's content is being sent.
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.
Result ID	No	No	This field indicates the Result ID of the result set of which this message's content is a part. This is normally used when a Producer is pushing Asynchronous Poll results (see Section 3.6.2.2).
Subscription Information	No	No	This field is only present if this message is being sent to provide content in fulfillment of an existing subscription. Absence of this field indicates that this message is not being sent in fulfillment of a subscription.
Collection Name [Feed Name]	Yes	No	This field indicates the name of the TAXII Data Collection from which this content is being provided.
Subscription ID	Yes	No	This field contains the Subscription ID for the subscription of which this content is being provided.

Name		Required?	Multiple?	Description
	Exclusive Begin Timestamp Label	No; At most one of these fields may appear	No	These fields serve the same purpose. Use of the Inclusive Begin Timestamp Label field is deprecated but retained for backwards compatibility with TAXII 1.0. Both fields MUST NOT appear together in the same message.
	Inclusive Begin Timestamp Label		No	Either field contains a Timestamp Label indicating the beginning of the time range this Inbox Message covers. (One field provides an exclusive value; the other provides an inclusive value.) Absence of either field indicates that the Inbox Message covers the earliest time for this TAXII Data Feed. The fields MUST NOT be included if the named TAXII Data Collection is a Data Set.
	Inclusive End Timestamp Label	Required if for a Feed; prohibited otherwise	No	This field contains a Timestamp Label indicating the end of the time range this Inbox Message covers. This field is inclusive. This field MUST be present if the named Data Collection is a Data Feed. This field MUST NOT be present if the named Data Collection is a Data Set.
Record Count		No	No	Indicates the number of applicable records for the given response, which MUST be greater than or equal to the number of content records returned in this message's Content Block(s). This field SHOULD be present in all Poll Response messages.
	Partial Count	No	No	This field indicates whether the provided Record Count is the exact number of applicable records, or if the provided number is a lower bound and there may be more records than stated. The field contains a boolean value. A value of TRUE indicates that the actual number of matching records may be greater than the value that appears in the Record Count field. A value of FALSE indicates the Record Count is an exact count of applicable records. If this field is absent, treat this field as having a value of FALSE.
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 4.5 for the definition of a Content Block.

The Destination Collection Name allows the sender of an Inbox Message to indicate one or more Data Collections to which the sender requests to have the enclosed content added. This can be used in a range of sharing models. The message recipient has full discretion as to whether to actually add the content to the indicated Data Collections as requested. For more details on the use of Destination Collection Name field and its use, see Section 3.2.1.

As with the Poll Response Message, the Inbox Message has two fields for beginning timestamp label values: the recommended Exclusive Begin Timestamp Label field and the deprecated Inclusive Begin Timestamp Label field. TAXII 1.1.1 implementations SHOULD use the Exclusive Begin Timestamp Label field if possible, but SHOULD support both for backwards compatibility. The use of these fields in the Inbox Message is identical to the use of these fields in the Poll Response Message.

4.4.11 TAXII[™] Poll Fulfillment Request

The TAXII Poll Fulfillment Request is used to collect results from a Poll Service where the result set has already been created. In general, this is used to collect results using Asynchronous Polling (see Section 3.6.2) or to collect multiple parts of a large result set over a Multi-Part Poll Exchange (see Section 3.6.1).

Name	Required?	Multiple?	Description
Collection Name	Yes	No	This field identifies the name of the TAXII Data Collection to which the request applies.
Result ID	Yes	No	The ID of the requested result set.
Result Part Number	Yes	No	If present, indicates the Result Part that is being collected.

4.5 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 4.1.7) or nesting expression (defined in Section 5.3) indicating the type of content contained in the Content field of this Content Block.
Subtype	No	No	This field identifies content binding subtypes of the specified Content Binding. Each Subtype MUST be a Content Binding Subtype ID as defined in the TAXII Content Binding Reference or by a third party. Absence of this field indicates that the content is not necessarily of any particular subtype.
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 a Timestamp Label associated with this Content Block. This field is only relevant if the content came from a TAXII Data Feed. It is at the sender's discretion as to whether this is included.
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.
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.

5 TAXII[™] 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 Collections in a TAXII Collection 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 for 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 Collection, 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 Data Collections and Content

TAXII Data Collections are how Producers expose content within a TAXII Architecture. Producers are allowed to assign content to Data Collections however they wish - Collections can align with 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 Collections.

5.2.1 TAXII[™] is Content Agnostic

The TAXII specifications do not provide details about the underlying content formats of records within TAXII. All content formats are a "black-box" as far as TAXII is concerned - none of the behaviors required to process TAXII at the message level 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 Data Feeds and Data Sets

TAXII supports two types of Data Collections: ordered Collections (Data Feeds) and unordered Collections (Data Sets). Data Producers can use either or both of these types of Collections as desired. The following sections look at the impact of these two types of Data Collections on TAXII activities.

5.2.2.1 Subscriptions

Subscriptions to Data Feeds and Data Sets behave in almost the same way, but there is one important difference. In the case of both Data Feeds and Data Sets, content **SHOULD** be considered for delivery when it is added to the Data Collection. From this initial set of considered content, eliminate records that do not match any provided Query expression, records that are expressed using a Content Binding that the subscriber doesn't support, or which are to be hidden from the subscriber under access control policies. The remaining records would then be delivered to the subscriber in accordance with the mechanisms outlined in the subscription request.

In addition, within a Data Set, any change to a record that was previously added to the Data Set **SHOULD** also be included for consideration. There is nothing in the TAXII Message to indicate that a given record is a revision of an older record rather than a new record that was just added to the Data Collection. Content might contain record identifiers that could be used to detect revision, but this is beyond the scope of TAXII.

5.2.2.2 Polling Ranges

When issuing a Poll Request, it is possible to limit the range of the request within a Data Feed. This is done by providing Timestamp Labels that serve as the bounds of the range to consider. Since Data Sets are not ordered, the same sort of restriction is not possible and all Poll Requests against Data Sets consider all records. (For this reason, a Consumer issuing a Poll Request against a Data Collection may wish to include a Query expression to limit the volume of material returned.)

When issuing a Poll Request against a Data Feed, Timestamp Labels play an important role in bounding the request. As noted in the section on TAXII Poll Request Messages (Section 4.4.8), Poll Requests that target a Data Feed 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 within a Data Feed, 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 result set if and only if the Producer is willing to share that content with the Consumer, the content matches any provided Query expression, 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 4.4.9, Producer MUST always include an upper bound in their Poll Response message when responding to a request against a Data Feed, 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 Data 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.

5.2.2.3 Data Feeds and Multi-Part Results

As described in Section 3.6.1, a very large result set may be broken up into parts for delivery. When the result set is created from a Data Feed, there is one additional rule that needs to be followed: each part of the result set MUST represent a contiguous range of the Data Feed and MUST contain all of the records that are part of the result set that fall within that range. This is required whether the results are being provided via Consumer-initiated pulls (section 3.6.1) or if they are being pushed to the Consumer, as can happen for very large result sets produced under Asynchronous Polling (see section 3.6.2.2). This is also required whether or not a Query expression was utilized in the creation of the result set. In other words, each individual message containing part of a result set needs to conform to the rules for Polling Ranges (Section 5.2.2.2). In effect, this requires that the entire result set for a Data Feed be sorted according to Timestamp Labels and then only broken apart so that each part represents exactly one segment of that sorted result set.

Data Sets, as they are unordered to begin with, do not have this requirement and any record could be sent in any part.

5.2.2.4 Content is Immutable within a Data Feed

As noted earlier in this document, within a Data Feed, 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 cannot be performed "in place" and would need to be accomplished by adding a new record (with a new, latest Timestamp Label) to the Data Feed. 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 a good way to indicate revision, correction, or revocation of that piece of content since 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.

By contrast, there is no requirement that Data Set content be immutable. Records within a Data Set may be modified or deleted freely. Note that Consumers are still on their own to determine if and when a particular record has been replaced, although, as noted in Section 5.2.2.1, Consumers subscribed to a Data Collection SHOULD be receive the modified record as part of their subscription.

5.2.3 Directing Inbox Message Content to Data Collections

As described in Section 3.2.1, TAXII supports the use of the Inbox Exchange to have pushed content added to one or more Data Collections. Some forms of hub-and-spoke architectures might make use of this feature if the hub is serving to resend submitted content on behalf of the spokes. Section 3.2.1 describes the exchanges that facilitate this behavior. A few additional notes about this capability are provided here:

- The recipient always has discretion as to whether the submitted content is added to a particular Data Collection. The recipient may reject content for any reason.

- It is also important to note that just because this specification refers to "automatically" adding content to Data Collections, TAXII does not require the act of adding content to be immediate. Recipients of content may wish for content to be reviewed, possibly by human operators, prior to its addition to a Data Collection.
- Content Blocks in Inbox Messages might include a Timestamp Label associated with any piece of submitted content. Even if the content is sent to a Data Feed, however, the recipient will assign the content's Timestamp Label when the content is added to the Feed. This may or may not be the same as the Timestamp Label provided when the content was sent.

In most cases, the individual agreements between sharing participants should stipulate if/how content sent to Data Collections should be handled. Specific questions that are not constrained by TAXII but which may need to be addressed in such agreements include:

- Can the recipient add the content to Data Collections other than the Data Collection(s) targeted by the sender?
- Can the recipient edit the content before it is added to the Data Collection(s)?
- Alternatively, is the recipient *required* to edit the content before it is added to Data Collections (for example, to anonymize the content prior to re-sharing)?
- Is only the content added to the Data Collection(s) or is the whole Content Block is added? (The latter might be desired as it would retain the original sender's signature, providing provenance, and any message associated with that content.)
- Whether certain access restrictions or other protections (e.g., encryption) will be employed if the content is re-shared.

5.2.4 Receive-Only Data Collections

Throughout this document Data Collections have been described almost exclusively with regard to Consumers using them to collect information from a Producer. It is important to note, however, that the creation of a Data Collection does not require that the contents of that Collection be made available to any Consumer. For example, Data Collections might be created to serve as "drop-boxes" for content sent from other parties. Inboxes could be configured to add submitted content to particular Data Collections, but those Data Collections might only be used internally by the recipient without ever making them available to any external Consumer. For this reason, it might be the case that a Collection Information Response Message includes a record for a Data Collection that includes neither a way to poll for its content nor any way to subscribe to it.

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 likely 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.

Content Binding Subtypes MUST NOT be used in an explicit nesting expression.

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 7 - 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 Collection, 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 Targeting Content Requests

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. Likewise, if the Consumer indicated support for a particular Content Binding Subtype of a given Content Binding, Producers **SHOULD NOT** send content expressed using an unsupported Content Binding Subtype.

Similarly, a subscription or a Poll Request may include a Query expression. Producers **SHOULD** not send records that do not result in a match against that Query expression. 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 sent by other parties.) If the request the Producer is fulfilling contains a Query, the Producer

SHOULD only send records for which it can determine there is a match. If the Producer has no access to the information needed to evaluate the query, this would result in no matching records.

On the other hand, if there is no Query expression, the Producer SHOULD send any content that might be acceptable to a Consumer based on supported Content types. 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.)

5.4.2 Paused Subscriptions

Subscribers to a Data Collection may request that delivery of content in fulfillment of a subscription be paused and later resumed. While a subscription is paused, the party sending content in fulfillment of the subscription MUST NOT send Inbox Messages to the subscriber in fulfillment of that subscription. When the subscriber resumes delivery, the party sending content in fulfillment of the subscription SHOULD deliver all content that would normally have been sent over the period the subscription was paused. Poll Servers MAY continue to serve Poll Requests associated with the subscription even while that subscription is in a paused state.

5.5 Query

Query expressions allow a Consumer to describe characteristics of interest within content records and limit the information that is collected from a Producer to records that have these characteristics. Query expressions can be added to Poll Requests or subscription management requests (i.e., Manage Collection Request Messages with an Action value of SUBSCRIBE). In the former case, they only constrain the results returned from that particular Poll Request. In the latter case, content will only be provided in fulfillment of the subscription if it matches against the provided Query expression.

Support of Query within TAXII is optional for both Consumers and Producers. Moreover, the TAXII Query capability is extensible, allowing Producers and Consumers to support different Query expression formats. Producers indicate which (if any) Query formats they support by listing the formats supported by particular TAXII Services within a Discovery Response message. A Poll Services MUST support one or more Query formats for Query expressions to be used on Poll Request Messages it receives. Likewise, Collection Management Service MUST support one or more Query formats for Query expressions to be used in subscription management requests. A Producer may choose to support Query expressions in only some of their services, and might choose to support different Query formats within different services.

TAXII defines a default Query format in the Default TAXII Query Specification. In addition, third parties may define their own Query formats for use in TAXII Query expressions. This section looks at how a Query format must be defined to be used within TAXII and how TAXII uses Query expressions to identify content of interest to a Consumer.

5.5.1 Query Format Specification Requirements

A Query Format Specification defines the rules surrounding the use of a particular Query format within TAXII. Each Query Specification MUST define a Query Format ID. The Query Format ID is used to identify the Query format within TAXII Messages. Query Format IDs must be globally unique. See Section 4.1.6 for more on Query Format IDs.

The Query field in both the Poll Request and the Manage Collection Subscription Request Messages only specifies a single sub-field, which contains the Query Format ID. All other information in the Query field, including any subfields and the type of information they need to contain, are left to the Query Format Specifications to define. The Query Format Specifications need to identify all sub-fields, describe the syntax of these fields (e.g., required or optional, etc.), how data is represented within them, and similar information. Moreover, it needs to describe how these fields are used to describe a set of criteria against which content records are compared. In particular, each Query Format Specification describes how to use

these fields and their contents to determine whether any given record is a "match" for a given Query expression.

Some Query formats may themselves have optional characteristics that may not be supported equally by all implementations. If this is the case, the Query Format Specification also needs to define one or more Supported Query subfields. The Supported Query subfields appear under the Supported Query field in the Discovery Response Message. (The Supported Query field already contains the Query Format ID subfield used to identify the Query format. However, like the Query field, other subfields are permitted.) The Query Format Specification can define subfields of Supported Query to indicate various parameters of a particular TAXII Service's support for a given Query Format. The Query Format Specification identifies the relevant subfields, describes their possible values, and dictates how they are to be interpreted by a Consumer in order to provide the necessary information to compose a Query expression that can be understood by that TAXII service.

In addition, because Query expressions and Supported Query subfields are conveyed in TAXII Messages that are expressed using some TAXII Message Binding, it is necessary to define Message Bindings for both the Query expression and Supported Query subfields for all Message Bindings over which this information is to be conveyed. These Message Bindings can be part of the Query Format Specification or part of some separate document.

5.5.2 General Query Processing

Most important aspects of processing Query expressions (namely, how to determine whether a record matches a given Query expression) are covered in the Query Format Specification. At the highest level, however, the processing of Query expressions is the same across all Query formats.

When a TAXII Service receives a message containing a Query expression, it should examine the Query Format ID to determine if a recognized Query format is present. If the Query Format is not supported, the TAXII Daemon SHOULD respond with a Status Message with a Status Type of "Unsupported Query Format". Otherwise, the Query expression is applied to the requested action. In the case of a Poll Request message, a result set is generated that only consists of records that match the Query expression. In the case of a subscription established with a Query expression, any record that would normally be sent in fulfillment of the subscription (new records in the case of a Data Feed; new or changed records in the case of a Data Set) is evaluated against the Query expression and only those that match the Query are sent to the subscriber in fulfillment of that particular subscription.

6 Conformance

Implementations have discretion over which parts of TAXII they implement (e.g., Discovery Service).

Conformant implementations must conform to all Normative Statements that apply to the portions of TAXII they implement (e.g., Implementers of the Discovery Service must conform to all Normative Statements regarding the Discovery Service).

Conformant implementations are free to ignore Normative Statements that do not apply to the portions of TAXII they implement (e.g., Non-implementers of the Discovery Service are free to ignore all Normative Statements regarding the Discovery Service).

The conformance section of this document is intentionally broad and attempts to reiterate what already exists in this document. The TAXII 1.1 Specifications, which this specification is based on, did not have a conformance section. Instead, the TAXII 1.1 Specifications relied on normative text. TAXII 1.1.1 represents a minimal change from TAXII 1.1, and in that spirit no new requirements have been defined in this document.

Appendix A. Acknowledgements

The individuals listed in this section have participated in the creation of this specification and are gratefully acknowledged.

Authors of the initial MITRE TAXII Specifications:

Mark Davidson, MITRE

Charles Schmidt, MITRE

Participants:

The following individuals were members of the OASIS CTI Technical Committee during the creation of this specification and their contributions are gratefully acknowledged:

- David Crawford, Aetna
- Joerg Eschweiler, Airbus Group SAS
- Marcos Orallo, Airbus Group SAS
- Roman Fiedler, AIT Austrian Institute of Technology
- Florian Skopik, AIT Austrian Institute of Technology
- Dean Thompson, Australia and New Zealand Banking Group (ANZ Bank)
- Alexander Foley, Bank of America
- Yogesh Mudgal, Bloomberg
- Owen Johnson, Blue Coat Systems, Inc.
- Bret Jordan, Blue Coat Systems, Inc.
- Adnan Baykal, Center for Internet Security (CIS)
- Ron Davidson, Check Point Software Technologies
- David McGrew, Cisco Systems
- Pavan Reddy, Cisco Systems
- Omar Santos, Cisco Systems
- Jyoti Verma, Cisco Systems
- Liron Schiff, Comilion (mobile) Ltd.
- Guy Wertheim, Comilion (mobile) Ltd.
- Doug DePeppe, Cyber Threat Intelligence Network, Inc. (CTIN)
- Jane Ginn, Cyber Threat Intelligence Network, Inc. (CTIN)
- Ben Othman, Cyber Threat Intelligence Network, Inc. (CTIN)
- Jeff Williams, Dell
- Richard Struse, DHS Office of Cybersecurity and Communications (CS&C)
- Marlon Taylor, DHS Office of Cybersecurity and Communications (CS&C)
- Dan Brown, DTCC
- Gordon Hundley, DTCC
- Chris Koutras, DTCC
- Robert Griffin, EMC
- Jeff Odom, EMC
- Ravi Sharda, EMC
- David Eilken, Financial Services Information Sharing and Analysis Center (FS-ISAC)
- Sarah Brown, Fox-IT
- Ryusuke Masuoka, Fujitsu Limited
- Eric Burger, Georgetown University
- Peter Allor, IBM
- Eldan Ben-Haim, IBM
- Peter Clark, IBM
- Sandra Hernandez, IBM

- Jason Keirstead, IBM
- John Morris, IBM
- Arvid Van Essche, IBM
- Ron Williams, IBM
- Paul Martini, iboss, Inc.
- Chris Richardson, IID
- Jerome Athias, Individual
- Peter Brown, Individual
- Elysa Jones, Individual
- Sanjiv Kalkar, Individual
- Bar Lockwood, Individual
- Terry MacDonald, Individual
- Alex Pinto, Individual
- Michael Schwartz, Individual
- Patrick Maroney, Integrated Networking Technologies, Inc.
- Andres More, Intel Corporation
- Wouter Bolsterlee, Intelworks BV
- Marko Dragoljevic, Intelworks BV
- Joep Gommers, Intelworks BV
- Sergey Polzunov, Intelworks BV
- Rutger Prins, Intelworks BV
- Andrei Sîrghi, Intelworks BV
- Raymon van der Velde, Intelworks BV
- Niels van Dijk, Intelworks BV
- Robert Huber, iSIGHT Partners, Inc.
- Ben Huguenin, Johns Hopkins University Applied Physics Laboratory
- Mark Moss, Johns Hopkins University Applied Physics Laboratory
- Pamela Smith, Johns Hopkins University Applied Physics Laboratory
- Terrence Driscoll, JPMorgan Chase Bank, N.A.
- David Laurance, JPMorgan Chase Bank, N.A.
- Brandon Hoffman, Lumeta Corporation
- Jonathan Baker, Mitre Corporation
- Sean Barnum, Mitre Corporation
- Mark Davidson, Mitre Corporation
- Jasen Jacobsen, Mitre Corporation
- Ivan Kirillov, Mitre Corporation
- Jon Salwen, Mitre Corporation
- Charles Schmidt, Mitre Corporation
- John Wunder, Mitre Corporation
- James Cabral, MTG Management Consultants, LLC.
- Scott Algeier, National Council of ISACs (NCI)
- Denise Anderson, National Council of ISACs (NCI)
- Josh Poster, National Council of ISACs (NCI)
- Mike Boyle, National Security Agency
- Jessica Fitzgerald-McKay, National Security Agency
- Takahiro Kakumaru, NEC Corporation
- John-Mark Gurney, New Context Services, Inc.
- Christian Hunt, New Context Services, Inc.
- Daniel Riedel, New Context Services, Inc.
- Andrew Storms, New Context Services, Inc.
- Nat Sakimura, Nomura Research Institute, Ltd. (NRI)
- David Darnell, North American Energy Standards Board
- Cory Casanave, Object Management Group
- Don Thibeu, Open Identity Exchange

- Vishaal Hariprasad, Palo Alto Networks
- John Tolbert, Queralto, Inc.
- Daniel Wyschogrod, Raytheon Company-SAS
- Ted Julian, Resilient Systems, Inc..
- Brian Engle, Retail Cyber Intelligence Sharing Center (R-CISC)
- Igor Baikarov, Securonix
- Bernd Grobauer, Siemens AG
- John Anderson, Soltra
- Aishwarya Asok Kumar, Soltra
- Peter Ayasse, Soltra
- Jeff Beekman, Soltra
- Jonathan Bush, Soltra
- Michael Butt, Soltra
- Cynthia Camacho, Soltra
- Aharon Chernin, Soltra
- Mark Clancy, Soltra
- Brady Cotton, Soltra
- Trey Darley, Soltra
- Paul Dion, Soltra
- Daniel Dye, Soltra
- Brandon Hanes, Soltra
- Robert Hutto, Soltra
- Ali Khan, Soltra
- Chris Kiehl, Soltra
- Michael Pepin, Soltra
- Natalie Suarez, Soltra
- David Waters, Soltra
- Chip Wickenden, Soltra
- Benjamin Yates, Soltra
- Cedric LeRoux, Splunk Inc.
- Brian Luger, Splunk Inc.
- Kathy Wang, Splunk Inc.
- Curtis Kostrosky, Symantec Corp.
- Greg Reaume, TELUS
- Alan Steer, TELUS
- Crystal Hayes, The Boeing Company
- Tyron Miller, Threat Intelligence Pty Ltd
- Andrew van der Stock, Threat Intelligence Pty Ltd
- Andrew Pendergast, ThreatConnect, Inc.
- Jason Spies, ThreatConnect, Inc.
- Nick Keuning, ThreatQuotient, Inc.
- Wei Huang, ThreatStream
- Hugh Njemanze, ThreatStream
- Chris Roblee, TruSTAR Technology
- Mark Angel, U.S. Bank
- Brad Butts, U.S. Bank
- Mona Magathan, U.S. Bank
- Adam Cooper, United Kingdom Cabinet Office
- Mike McLellan, United Kingdom Cabinet Office
- Chris O'Brien, United Kingdom Cabinet Office
- James Penman, United Kingdom Cabinet Office
- Howard Staple, United Kingdom Cabinet Office
- Alastair Treharne, United Kingdom Cabinet Office
- Julian White, United Kingdom Cabinet Office

- Evette Maynard-Noel, US Department of Homeland Security
 - Justin Stekervetz, US Department of Homeland Security
 - Robert Coderre, VeriSign
 - Kyle Maxwell, VeriSign
 - Lee Chieffalo, ViaSat, Inc.
 - Wilson Figueroa, ViaSat, Inc.
 - Anthony Rutkowski, Yaana Technologies, LLC
-
-

Special Thanks:

A special thanks to the US Department of Homeland Security's (DHS) National Cybersecurity and Communications Integration Center (NCCIC), and to Richard Struse, Chief Advanced Technology Officer of the DHS NCCIC. Without your sponsorship, vision, and relentless vigor, none of this would have been possible.

Appendix B. Revision History

Revision	Date	Editor	Changes Made
Working Draft 01	July 2015	Mark Davidson	Initial working draft based on MITRE TAXII Services Specification 1.1.