Internet Engineering Task Force (IETF)
Request for Comments: 6759
Category: Informational
ISSN: 2070-1721

B. Claise
P. Aitken
N. Ben-Dvora
Cisco Systems, Inc.

November 2012

Cisco Systems Export of Application Information in IP Flow Information Export (IPFIX)

Abstract

This document specifies a Cisco Systems extension to the IPFIX information model specified in RFC 5102 to export application information.

Status of This Memo

Copyright Notice

Table of Contents

List of Figures

List of Tables

1. Introduction

Today, service providers and network administrators are looking for visibility into the packet content rather than just the packet header. Some network devices' Metering Processes inspect the packet content and identify the applications that are utilizing the network traffic. Applications in this context are defined as networking protocols used by networking processes that exchange packets between them (such as web applications, peer-to-peer applications, file transfer, e-mail applications, etc.). Applications can be further characterized by other criteria, some of which are application specific. Examples include: web application to a specific domain, per-user specific traffic, a video application with a specific codec, etc.

The application identification is based on several different methods or even a combination of methods:

1. L2 (Layer 2) protocols (such as ARP (Address Resolution Protocol), PPP (Point-to-Point Protocol), LLDP (Link Layer Discovery Protocol))

2. IP protocols (such as ICMP (Internet Control Message Protocol), IGMP (Internet Group Management Protocol), GRE (Generic Routing Encapsulation)

3. TCP or UDP ports (such as HTTP, Telnet, FTP)

4. Application layer header (of the application to be identified)

5. Packet data content

6. Packets and traffic behavior

The exact application identification methods are part of the Metering Process internals that aim to provide an accurate identification and minimize false identification. This task requires a sophisticated Metering Process since the protocols do not behave in a standard manner.

1. Applications use port obfuscation where the application runs on a different port than the IANA assigned one. For example, an HTTP server might run on TCP port 23 (assigned to telnet in [IANA-PORTS]).

2. IANA port registries do not accurately reflect how certain ports are "commonly" used today. Some ports are reserved, but the application either never became prevalent or is not in use today.

3. The application behavior and identification logic become more and more complex.

For that reason, such Metering Processes usually detect applications based on multiple mechanisms in parallel. Detection based only on port matching might wrongly identify the application. If the Metering Process is capable of detecting applications more accurately, it is considered to be stronger and more accurate.

Similarly, a reporting mechanism that uses L4 port based applications only, such as L4:<known port>, would have similar issues. The reporting system should be capable of reporting the applications classified using all types of mechanisms. In particular, applications that do not have any IANA port definition. While a mechanism to export application information should be defined, the L4 port being used must be exported using the destination port (destinationTransportPort at [IANA-IPFIX]) in the corresponding IPFIX record.

Applications could be identified at different OSI layers, from layer 2 to layer 7. For example, the Link Layer Distribution Protocol (LLDP) [LLDP] can be identified in layer 2, ICMP can be identified in layer 3 [IANA-PROTO], HTTP can be identified in layer 4 [IANA-PORTS], and Webex can be identified in layer 7.

While an ideal solution would be an IANA registry for applications above (or inside the payload of) the well-known ports [IANA-PORTS], this solution is not always possible. Indeed, the specifications for some applications embedded in the payload are not available. Some reverse engineering as well as a ubiquitous language for application identification would be required conditions to be able to manage an IANA registry for these types of applications. Clearly, these are blocking factors.

This document specifies the Cisco Systems application information encoding (as described in Section 4) to export the application information with the IPFIX protocol [RFC5101]. However, the layer 7 application registry values are out of scope of this document.

1.1. Application Information Use Cases

There are several use cases for application information:

1. Application Visibility

   This is one of the main cases for using application information. Network administrators are using application visibility to understand the main network consumers, network trends, and user behavior.

2. Security Functions

   Application knowledge is sometimes used in security functions in order to provide comprehensive functions such as Application-based firewall, URL filtering, parental control, intrusion detection, etc.

All of the above use cases require exporting application information to provide the network function itself or to log the network function operation.

1.2. Conventions Used in This Document

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 RFC 2119 [RFC2119].

2. IPFIX Documents Overview

The IPFIX protocol [RFC5101] provides network administrators with access to IP Flow information.

The architecture for the export of measured IP Flow information out of an IPFIX Exporting Process to a Collecting Process is defined in the IPFIX Architecture [RFC5470], per the requirements defined in RFC 3917 [RFC3917].

The IPFIX Architecture [RFC5470] specifies how IPFIX Data Records and Templates are carried via a congestion-aware transport protocol from IPFIX Exporting Processes to IPFIX Collecting Processes.

IPFIX has a formal description of IPFIX Information Elements, their names, types, and additional semantic information, as specified in the IPFIX information model [RFC5102].

In order to gain a level of confidence in the IPFIX implementation, probe the conformity and robustness, and allow interoperability, the Guidelines for IPFIX Testing [RFC5471] presents a list of tests for implementers of compliant Exporting Processes and Collecting Processes.

The Bidirectional Flow Export [RFC5103] specifies a method for exporting bidirectional flow (biflow) information using the IPFIX protocol, representing each biflow using a single Flow Record.

"Reducing Redundancy in IP Flow Information Export (IPFIX) and Packet Sampling (PSAMP) Reports" [RFC5473] specifies a bandwidth-saving method for exporting Flow or packet information, by separating information common to several Flow Records from information specific to an individual Flow Record: common Flow information is exported only once.

3. Terminology

IPFIX-specific terminology used in this document is defined in Section 2 of the IPFIX protocol specification [RFC5101]. As in [RFC5101], these IPFIX-specific terms have the first letter of a word capitalized when used in this document.

3.1. New Terminology

Application ID

A unique identifier for an application.

When an application is detected, the most granular application is encoded in the Application ID.

4. applicationId Information Element Specification

This document specifies the applicationId Information Element, which is a single field composed of two parts:

1. 8 bits of Classification Engine ID. The Classification Engine can be considered as a specific registry for application assignments.

2. n bits of Selector ID. The Selector ID length varies depending on the Classification Engine ID.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Class. Eng. ID|         Selector ID  ...                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              ...                              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 1: applicationId Information Element

Classification Engine ID

A unique identifier for the engine that determined the Selector ID. Thus, the Classification Engine ID defines the context for the Selector ID.

Selector ID

A unique identifier of the application for a specific Classification Engine ID. Note that the Selector ID length varies depending on the Classification Engine ID.

The Selector ID term is a similar concept to the selectorId Information Element, specified in the PSAMP Protocol [RFC5476][RFC5477].

4.1. Existing Classification Engine IDs

The following Classification Engine IDs have been allocated:

Name	Value	Description
	0	Invalid.
IANA-L3	1	The Assigned Internet Protocol Number (layer 3 (L3)) is exported in the Selector ID. See [IANA-PROTO].
PANA-L3	2	Proprietary layer 3 definition. An enterprise can export its own layer 3 protocol numbers. The Selector ID has a global significance for all devices from the same enterprise.

IANA-L4	3	The IANA layer 4 (L4) well-known port number is exported in the Selector ID. See [IANA-PORTS]. Note: as an IPFIX flow is unidirectional, it contains the destination port.
PANA-L4	4	Proprietary layer 4 definition. An enterprise can export its own layer 4 port numbers. The Selector ID has global significance for devices from the same enterprise. Example: IPFIX was pre-assigned the port 4739 using the IANA early allocation process [RFC4020] years before the document was published as an RFC. While waiting for the RFC and its associated IANA registration, Selector ID 4739 was used with this PANA-L4.
	5	Reserved.
USER-Defined	6	The Selector ID represents applications defined by the user (using CLI, GUI, etc.) based on the methods described in Section 1. The Selector ID has a local significance per device.
	7	Reserved.
	8	Reserved.
	9	Reserved.
	10	Reserved.
	11	Reserved.
PANA-L2	12	Proprietary layer 2 (L2) definition. An enterprise can export its own layer 2 identifiers. The Selector ID represents the enterprise's unique global layer 2 applications. The Selector ID has a global significance for all

		devices from the same enterprise. Examples include Cisco Subnetwork Access Protocol (SNAP).
PANA-L7	13	Proprietary layer 7 definition. The Selector ID represents the enterprise's unique global ID for layer 7 applications. The Selector ID has a global significance for all devices from the same enterprise. This Classification Engine ID is used when the application registry is owned by the Exporter manufacturer (referred to as the "enterprise" in this document).
	14	Reserved.
	15	Reserved.
	16	Reserved.
	17	Reserved.
ETHERTYPE	18	The Selector ID represents the well-known Ethertype. See [ETHERTYPE].
LLC	19	The Selector ID represents the well-known IEEE 802.2 Link Layer Control (LLC) Destination Service Access Point (DSAP). See [LLC].
PANA-L7-PEN	20	Proprietary layer 7 definition, including a Private Enterprise Number (PEN) [IANA-PEN] to identify that the application registry being used is not owned by the Exporter manufacturer or to identify the original enterprise in the case of a mediator or 3rd party device. The Selector ID represents the enterprise unique global ID for the layer 7 applications. The

	Selector ID has a global significance for all devices from the same enterprise.
21 to 255	Available (255 is the maximum Engine ID)

Table 1: Existing Classification Engine IDs

"PANA = Proprietary Assigned Number Authority". In other words, an enterprise specific version of IANA for internal IDs.

The PANA-L7 Classification Engine ID SHOULD be used when the application registry is owned by the Exporter manufacturer. Even if the application registry is owned by the Exporter manufacturer, the PANA-L7-PEN MAY be used, specifying the manufacturer.

For example, if Exporter A (from enterprise-A) wants to export its enterprise-A L7 registry, then it uses the PANA-L7 Classification Engine ID. If Exporter B (from enterprise-B) wants to export its enterprise-B L7 registry, then it also uses the PANA-L7 Classification Engine ID.

The mechanism for the Collector to know about the Exporter PEN is out of scope of this document. Possible tracks are SNMP polling, an Options Template exporting the privateEnterpriseNumber Information Element [IANA-IPFIX], hardcoded value, etc.

An Exporter may classify the application according to another vendor's application registry. For example, an IPFIX Mediator [RFC6183] may need to re-export applications received from different Exporters using different PANA-L7 application registries. For example, if Exporter C (from enterprise-C) wants to reuse enterprise-D's application registry, then it uses PANA-L7-PEN with enterprise-D's PEN.

When reporting application information from multiple Exporters from different enterprises (different PENs), the PANA-L7-PEN Classification Engine MUST be used in exported Flow Records, which allows the original enterprise ID to be reported. The ID of the enterprise that defined the Application ID is identified by the enterprise's PEN. For example, an IPFIX Mediator aggregates traffic from some Exporters which report enterprise-E applications and other Exporters that report enterprise-F applications.

An example is displayed in Section 6.6.

Note that the PANA-L7 Classification Engine ID is also used for resolving IANA L4 port Discrepancies (see Section 4.4).

The list in Table 1 is maintained by IANA thanks to the registry within the classificationEngineId Information Element. See the IANA Considerations section. The Classification Engine ID is part of the Application ID encoding, so the classificationEngineId Information Element is currently not required by the specifications in this document. However, this Information Element was created for completeness, as it was anticipated that this Information Element will be required in the future.

4.2. Selector ID Length per Classification ID

As the Selector ID part of the Application ID is variable based on the Classification Engine ID value, the applicationId SHOULD be encoded in a variable-length Information Element [RFC5101] for IPFIX export.

The following table displays the Selector ID default length for the different Classification Engine IDs.

Classification Engine ID Name	Selector ID default length (in bytes)
IANA-L3	1
PANA-L3	1
IANA-L4	2
PANA-L4	2
USER-Defined	3
PANA-L2	5
PANA-L7	3
ETHERTYPE	2
LLC	1
PANA-L7-PEN	3 (*)

Table 2: Selector ID Default Length per Classification Engine ID

(*) There are an extra 4 bytes for the PEN. However, the PEN is not considered part of the Selector ID.

If a legacy protocol such as NetFlow version 9 [RFC3954] is used, and this protocol doesn't support variable-length Information Elements, then either multiple Template Records (one per applicationId length), or a single Template Record corresponding to the maximum sized applicationId MUST be used.

Application IDs MAY be encoded in a smaller number of bytes, following the same rules as for IPFIX Reduced Size Encoding [RFC5101].

Application IDs MAY be encoded with a larger length. For example, a normal IANA L3 protocol encoding would take 2 bytes since the Selector ID represents the protocol field from the IP header encoded in one byte. However, an IANA L3 protocol encoding may be encoded with 3 bytes. In this case, the Selector ID value MUST always be encoded in the least significant bits as shown in Figure 2.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|Class. Eng. ID |zero-valued upper-bits ... Selector ID         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 2: Selector ID Encoding

4.3. Application Name Options Template Record

For Classification Engines that specify locally unique Application IDs (which means unique per engine and per router), an Options Template Record (see [RFC5101]) MUST be used to export the correspondence between the Application ID, the Application Name, and the Application Description.

For Classification Engines that specify globally unique Application IDs, an Options Template Record MAY be used to export the correspondence between the Application ID, the Application Name and the Application Description, unless the mapping is hardcoded in the Collector, or known out of band (for example, by polling a MIB).

An example Options Template is shown in Section 6.8.

Enterprises may assign company-wide Application ID values for the PANA-L7 Classification Engine. In this case, a possible optimization for the Collector is to keep the mappings between the Application IDs and the Application Names per enterprise, as opposed to per Exporter.

4.4. Resolving IANA L4 Port Discrepancies

Even though IANA L4 ports usually point to the same protocols for both UDP, TCP or other transport types, there are some exceptions, as mentioned in Appendix B.

Instead of imposing the transport protocol (UDP/TCP/SCTP/etc.) in the scope of the "Application Name Options Template Record" (Section 6.8) for all applications (in addition to having the transport protocol as a key-field in the Flow Record definition), the convention is that the L4 application is always TCP related. So, whenever the Collector has a conflict in looking up IANA, it would choose the TCP choice. As a result, the UDP L4 applications from Table 3 and the SCTP L4 applications from Table 4 are assigned in the PANA_L7 Application ID range, i.e., under Classification Engine ID 13.

Currently, there are no discrepancies between the well-known ports for TCP and the Datagram Congestion Control Protocol (DCCP).

5. Grouping Applications with Attributes

Due to the high number of different Application IDs, Application IDs MAY be categorized into groups. This offers the benefits of easier reporting and action, such as QoS policies. Indeed, most applications with the same characteristics should be treated the same way; for example, all video traffic.

Attributes are statically assigned per Application ID and are independent of the traffic. The attributes are listed below:

Name	Description
Category	An attribute that provides a first-level categorization for each Application ID. Examples include browsing, email, file-sharing, gaming, instant messaging, voice-and-video, etc. The category attribute is encoded by the applicationCategoryName Information Element.
Sub-Category	An attribute that provides a second-level categorization for each Application ID. Examples include backup-systems, client-server, database, routing-protocol, etc. The sub-category attribute is

	encoded by the applicationSubCategoryName Information Element.
Application-Group	An attribute that groups multiple Application IDs that belong to the same networking application. For example, the ftp-group contains ftp-data (port 20), ftp (port 20), ni-ftp (port 47), sftp (port 115), bftp (port 152), ftp-agent(port 574), ftps-data (port 989). The application-group attribute is encoded by the applicationGroupName Information Element.
P2P-Technology	Specifies if the Application ID is based on peer-to-peer technology. The P2P-technology attribute is encoded by the p2pTechnology Information Element.
Tunnel-Technology	Specifies if the Application ID is used as a tunnel technology. The tunnel-technology attribute is encoded by the tunnelTechnology Information Element.
Encrypted	Specifies if the Application ID is an encrypted networking protocol. The encrypted attribute is encoded by the encryptedTechnology Information Element.

Table 3: Application ID Static Attributes

Every application is assigned to one applicationCategoryName, one applicationSubCategoryName, one applicationGroupName, and it has one p2pTechnology, one tunnelTechnology, and one encryptedTechnology. These new Information Elements are specified in the IANA Considerations section (Section 7.1).

Maintaining the attribute values in IANA seems impossible to realize. Therefore, the attribute values per application are enterprise specific.

5.1. Options Template Record for Attribute Values

An Options Template Record (see [RFC5101]) SHOULD be used to export the correspondence between each Application ID and its related Attribute values. An alternative way for the Collecting Process to learn the correspondence is to populate these mappings out of band, for example, by loading a CSV file containing the correspondence table.

The Attributes Option Template contains the application ID as a scope field, followed by the applicationCategoryName, the applicationSubCategoryName, the applicationGroupName, the p2pTechnology, the tunnelTechnology, and the encryptedTechnology Information Elements.

A list of attributes may conveniently be exported using a subTemplateList per [RFC6313].

An example is given in Section 6.9.

6. Application ID Examples

The following examples are created solely for the purpose of illustrating how the extensions proposed in this document are encoded.

6.1. Example 1: Layer 2 Protocol

The list of Classification Engine IDs in Table 1 shows that the layer 2 Classification Engine IDs are 12 (PANA-L2), 18, (ETHERTYPE) and 19 (LLC).

From the Ethertype list, LLDP [LLDP] has the Selector ID value 0x88CC, so 35020 in decimal:

NAME	Selector ID
LLDP	35020

So, in the case of LLDP, the Classification Engine ID is 18 (LLC) while the Selector ID has the value 35020.

Per Section 4, the applicationId Information Element is a single field composed of 8 bits of Classification Engine ID, followed by n bits of Selector ID. From Table 2, the Selector ID length n is 2 for the ETHERTYPE Engine ID.

Therefore, the Application ID is encoded as:

```
 0                   1                   2
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|       18      |             35020             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

So the Application ID has the decimal value of 1214668. The format '18..35020' is used for simplicity in the examples below, to clearly express that two components of the Application ID.

The Exporting Process creates a Template Record with a few Information Elements: amongst other things, the Application ID. For example:

- applicationId (key field)
- octetTotalCount (non-key field)

For example, a Flow Record corresponding to the above Template Record may contain:

```
{ applicationId='18..35020',
  octetTotalCount=123456 }
```

The Collector has all the required information to determine that the application is LLDP, because the Application ID uses a global and well-known registry, i.e., the Ethertype. The Collector can determine which application is represented by the Application ID by loading the registry out of band.

6.2. Example 2: Standardized IANA Layer 3 Protocol

From the list of Classification Engine IDs in Table 1, the IANA layer 3 Classification Engine ID (IANA-L3) is 1. From Table 2 the Selector ID length is 1 for the IANA-L3 Engine ID.

From the list of IANA layer 3 protocols (see [IANA-PROTO]), ICMP has the value 1:

Decimal	Keyword	Protocol	Reference
1	ICMP	Internet Control Message	[RFC792]

So, in the case of the standardized IANA layer 3 protocol ICMP, the Classification Engine ID is 1, and the Selector ID has the value of 1.

Therefore, the Application ID is encoded as:

```
     0                   1
     0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |       1       |       1       |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

So, the Application ID has the value of 257. The format '1..1' is used for simplicity in the examples below.

The Exporting Process creates a Template Record with a few Information Elements: amongst other things, the Application ID. For example:

- sourceIPv4Address (key field)
- destinationIPv4Address (key field)
- ipDiffServCodePoint (key field)
- applicationId (key field)
- octetTotalCount (non-key field)

For example, a Flow Record corresponding to the above Template Record may contain:

```
{ sourceIPv4Address=192.0.2.1,
  destinationIPv4Address=192.0.2.2,
  ipDiffServCodePoint=0,
  applicationId='1..1',
  octetTotalCount=123456 }
```

The Collector has all the required information to determine that the application is ICMP, because the Application ID uses a global and well-known registry, i.e., the IANA L3 protocol number.

6.3. Example 3: Proprietary Layer 3 Protocol

Assume that an enterprise has specified a new layer 3 protocol called "foo".

From the list of Classification Engine IDs in Table 1, the proprietary layer 3 Classification Engine ID (PANA-L3) is 2. From Table 2 the Selector ID length is 1 for the PANA-L3 Engine ID.

A global registry within the enterprise specifies that the "foo" protocol has the value 90:

Protocol	Protocol ID
foo	90

So, in the case of the layer 3 protocol foo specified by this enterprise, the Classification Engine ID is 2, and the Selector ID has the value of 90.

Therefore, the Application ID is encoded as:

```
 0                   1
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|       2       |      90       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

So the Application ID has the value of 602. The format '2..90' is used for simplicity in the examples below.

The Exporting Process creates a Template Record with a few Information Elements: amongst other things, the Application ID. For example:

- sourceIPv4Address (key field)
- destinationIPv4Address (key field)
- ipDiffServCodePoint (key field)
- applicationId (key field)
- octetTotalCount (non-key field)

For example, a Flow Record corresponding to the above Template Record may contain:

```
{ sourceIPv4Address=192.0.2.1,
  destinationIPv4Address=192.0.2.2,
  ipDiffServCodePoint=0,
  applicationId='2..90',
  octetTotalCount=123456 }
```

Along with this Flow Record, a new Options Template Record would be exported, as shown in Section 6.8.

6.4. Example 4: Standardized IANA Layer 4 Port

From the list of Classification Engine IDs in Table 1, the IANA layer 4 Classification Engine ID (IANA-L4) is 3. From Table 2 the Selector ID length is 2 for the IANA-L4 Engine ID.

From the list of IANA layer 4 ports (see [IANA-PORTS]), SNMP has the value 161:

Keyword	Decimal	Description
snmp	161/tcp	SNMP
snmp	161/udp	SNMP

So, in the case of the standardized IANA layer 4 SNMP port, the Classification Engine ID is 3, and the Selector ID has the value of 161.

Therefore, the Application ID is encoded as:

```
 0                   1
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|       3       |              161              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

So the Application ID has the value of 196769. The format '3..161' is used for simplicity in the examples below.

The Exporting Process creates a Template Record with a few Information Elements: amongst other things, the Application ID. For example:

- sourceIPv4Address (key field)
- destinationIPv4Address (key field)
- protocol (key field)
- ipDiffServCodePoint (key field)
- applicationId (key field)
- octetTotalCount (non-key field)

For example, a Flow Record corresponding to the above Template Record may contain:

```
{ sourceIPv4Address=192.0.2.1,
  destinationIPv4Address=192.0.2.2,
  protocol=17, ipDiffServCodePoint=0,
  applicationId='3..161',
  octetTotalCount=123456 }
```

The Collector has all the required information to determine that the application is SNMP, because the Application ID uses a global and well-known registry, i.e., the IANA L4 protocol number.

6.5. Example 5: Layer 7 Application

In this example, the Metering Process has observed some Webex traffic.

From the list of Classification Engine IDs in Table 1, the layer 7 unique Classification Engine ID (PANA-L7) is 13. From Table 2 the Selector ID length is 3 for the PANA-L7 Engine ID.

Suppose that the Metering Process returns the ID 10000 for Webex traffic.

So, in the case of this Webex application, the Classification Engine ID is 13 and the Selector ID has the value of 10000.

Therefore, the Application ID is encoded as:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|      13       |                    10000                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

So the Application ID has the value of 218113808. The format '13..10000' is used for simplicity in the examples below.

The Exporting Process creates a Template Record with a few Information Elements: amongst other things, the Application ID. For example:

- sourceIPv4Address (key field)
- destinationIPv4Address (key field)
- ipDiffServCodePoint (key field)
- applicationId (key field)
- octetTotalCount (non-key field)

For example, a Flow Record corresponding to the above Template Record may contain:

```
{ sourceIPv4Address=192.0.2.1,
  destinationIPv4Address=192.0.2.2,
  ipDiffServCodePoint=0,
  applicationId='13..10000',
  octetTotalCount=123456 }
```

The 10000 value is globally unique for the enterprise, so that the Collector can determine which application is represented by the Application ID by loading the registry out of band.

Along with this Flow Record, a new Options Template Record would be exported, as shown in Section 6.8.

6.6. Example 6: Layer 7 Application with Private Enterprise Number (PEN)

In this example, the layer 7 Webex traffic from Example 5 above have been classified by enterprise X. The exported records have been received by enterprise Y's mediation device, which wishes to forward them to a top-level Collector.

In order for the top-level Collector to know that the records were classified by enterprise X, the enterprise Y mediation device must report the records using the PANA-L7-PEN Classification Engine ID with enterprise X's Private Enterprise Number.

The PANA-L7-PEN Classification Engine ID is 20, and enterprise X's Selector ID for Webex traffic has the value of 10000. From Table 2 the Selector ID length is 3 for the PANA-L7-PEN Engine ID.

Therefore, the Application ID is encoded as:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|      20       |               enterprise ID = X            ...|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|...Ent.ID.contd|                     10000                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The format '20..X..10000' is used for simplicity in the examples below.

The Exporting Process creates a Template Record with a few Information Elements: amongst other things, the Application ID. For example:

- sourceIPv4Address (key field)
- destinationIPv4Address (key field)
- ipDiffServCodePoint (key field)
- applicationId (key field)
- octetTotalCount (non-key field)

For example, a Flow Record corresponding to the above Template Record may contain:

```
{ sourceIPv4Address=192.0.2.1,
  destinationIPv4Address=192.0.2.2,
  ipDiffServCodePoint=0,
  applicationId='20..X..10000',
  octetTotalCount=123456 }
```

The 10000 value is globally unique for enterprise X, so that the Collector can determine which application is represented by the Application ID by loading the registry out of band.

Along with this Flow Record, a new Options Template Record would be exported, as shown in Section 6.8.

6.7. Example: Port Obfuscation

For example, an HTTP server might run on a TCP port 23 (assigned to telnet in [IANA-PORTS]). If the Metering Process is capable of detecting HTTP in the same case, the Application ID representation must contain HTTP. However, if the reporting application wants to determine whether or not the default HTTP port 80 or 8080 was used, the transport ports (sourceTransportPort and destinationTransportPort at [IANA-IPFIX]) must also be exported in the corresponding IPFIX record.

In the case of a standardized IANA layer 4 port, the Classification Engine ID (PANA-L4) is 3, and the Selector ID has the value of 80 for HTTP (see [IANA-PORTS]). From Table 2 the Selector ID length is 2 for the PANA-L4 Engine ID.

Therefore, the Application ID is encoded as:

```
 0                   1                   2
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|       3       |              80               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The Exporting Process creates a Template Record with a few Information Elements: amongst other things, the Application ID. For example:

- sourceIPv4Address (key field)
- destinationIPv4Address (key field)
- protocol (key field)
- destinationTransportPort (key field)
- applicationId (key field)
- octetTotalCount (non-key field)

For example, a Flow Record corresponding to the above Template Record may contain:

```
{ sourceIPv4Address=192.0.2.1,
  destinationIPv4Address=192.0.2.2,
  protocol=17,
  destinationTransportPort=23,
  applicationId='3..80',
  octetTotalCount=123456 }
```

The Collector has all the required information to determine that the application is HTTP, but runs on port 23.

6.8. Example: Application Name Mapping Options Template

Along with the Flow Records shown in the above examples, a new Options Template Record should be exported to express the Application Name and Application Description associated with each Application ID.

The Options Template Record contains the following Information Elements:

1. Scope = applicationId.

   From RFC 5101: The scope, which is only available in the Options Template Set, gives the context of the reported Information Elements in the Data Records.

2. applicationName.

3. applicationDescription.

The Options Data Record associated with the examples above would contain, for example:

```
{ scope=applicationId='2..90',
  applicationName="foo",
  applicationDescription="The foo protocol",

  scope=applicationId='13..10000',
  applicationName="webex",
  applicationDescription="Webex application" }

  scope=applicationId='20..X..10000',
  applicationName="webex",
  applicationDescription="Webex application" }
```

When combined with the example Flow Records above, these Options Template Records tell the Collector:

1. A flow of 123456 bytes exists from sourceIPv4Address 192.0.2.1 to destinationIPv4address 192.0.2.2 with an applicationId of '12..90', which maps to the "foo" application.

2. A flow of 123456 bytes exists from sourceIPv4Address 192.0.2.1 to destinationIPv4address 192.0.2.2 with an Application ID of '13..10000', which maps to the "Webex" application.

3. A flow of 123456 bytes exists from sourceIPv4Address 192.0.2.1 to destinationIPv4address 192.0.2.2 with an Application ID of '20..PEN..10000', which maps to the "Webex" application, according to the application registry from the enterprise X.

6.9. Example: Attributes Values Options Template Record

Along with the Flow Records shown in the above examples, a new Options Template Record is exported to express the values of the different attributes related to the Application IDs.

The Options Template Record would contain the following Information Elements:

1. Scope = applicationId.

   From RFC 5101: The scope, which is only available in the Options Template Set, gives the context of the reported Information Elements in the Data Records.

2. applicationCategoryName.

3. applicationSubCategoryName.

4. applicationGroupName

5. p2pTechnology

6. tunnelTechnology

7. encryptedTechnology

The Options Data Record associated with the examples above would contain, for example:

```
{ scope=applicationId='2..90',
  applicationCategoryName="foo-category",
  applicationSubCategoryName="foo-subcategory",
  applicationGroupName="foo-group",
  p2pTechnology=NO
  tunnelTechnology=YES
  encryptedTechnology=NO
```

When combined with the example Flow Records above, these Options Template Records tell the Collector:

A flow of 123456 bytes exists from sourceIPv4Address 192.0.2.1 to destinationIPv4address 192.0.2.2 with a DSCP value of 0 and an applicationId of '12..90', which maps to the "foo" application. This application can be characterized by the relevant attributes values.

7. IANA Considerations

7.1. New Information Elements

This document specifies 10 new IPFIX Information Elements: applicationDescription, applicationId, applicationName, classificationEngineId, applicationCategoryName, applicationSubCategoryName, applicationGroupName, p2pTechnology, tunnelTechnology, and encryptedTechnology.

The new Information Elements listed below have been added to the IPFIX Information Element registry at [IANA-IPFIX].

7.1.1. applicationDescription

Name: applicationDescription
Description:
  Specifies the description of an application.
Abstract Data Type: string
Data Type Semantics:
ElementId: 94
Status: current

7.1.2. applicationId

Name: applicationId
Description:
Specifies an Application ID.
Abstract Data Type: octetArray
Data Type Semantics: identifier
Reference: See Section 4 of [RFC6759]
for the applicationId Information Element Specification.
ElementId: 95
Status: current

7.1.3. applicationName

Name: applicationName
Description:
Specifies the name of an application.
Abstract Data Type: string
Data Type Semantics:
ElementId: 96
Status: current

7.1.4. classificationEngineId

Name: classificationEngineId
Description:
A unique identifier for the engine that determined the Selector ID. Thus, the Classification Engine ID defines the context for the Selector ID. The Classification Engine can be considered as a specific registry for application assignments.

Initial values for this field are listed below. Further values may be assigned by IANA in the Classification Engine IDs registry per Section 7.2.

0 Invalid.

1 IANA-L3: The Assigned Internet Protocol Number (layer 3 (L3)) is exported in the Selector ID. See http://www.iana.org/assignments/protocol-numbers.

2 PANA-L3: Proprietary layer 3 definition. An enterprise can export its own layer 3 protocol numbers. The Selector ID has a global significance for all devices from the same enterprise.

3 IANA-L4: The IANA layer 4 (L4) well-known port number is exported in the Selector ID. See [IANA-PORTS]. Note: as an IPFIX flow is unidirectional, it contains the destination port.

4 PANA-L4: Proprietary layer 4 definition. An enterprise can export its own layer 4 port numbers. The Selector ID has global significance for devices from the same enterprise. Example: IPFIX was pre-assigned port 4739 using the IANA early allocation process [RFC4020] years before the document was published as an RFC. While waiting for the RFC and it associated IANA registration, Selector ID 4739 was used with this PANA-L4.

5 Reserved

6 USER-Defined: The Selector ID represents applications defined by the user (using CLI, GUI, etc.) based on the methods described in Section 2. The Selector ID has a local significance per device.

7 Reserved

8 Reserved

9 Reserved

10 Reserved

11 Reserved

12 PANA-L2: Proprietary layer 2 (L2) definition. An enterprise can export its own layer 2 identifiers. The Selector ID represents the enterprise’s unique global layer 2 applications. The Selector ID has a global significance for all devices from the same enterprise. Examples include the Cisco Subnetwork Access Protocol (SNAP).

13 PANA-L7: Proprietary layer 7 definition. The Selector ID represents the enterprise's unique global ID for layer 7 applications. The Selector ID has a global significance for all devices from the same enterprise. This Classification Engine ID is used when the application registry is owned by the Exporter manufacturer (referred to as the "enterprise" in this document).

14 Reserved

15 Reserved

16 Reserved

17 Reserved

18 ETHERTYPE: The Selector ID represents the well-known Ethertype. See [ETHERTYPE].

19 LLC: The Selector ID represents the well-known IEEE 802.2 Link Layer Control (LLC) Destination Service Access Point (DSAP). See [LLC].

20 PANA-L7-PEN: Proprietary layer 7 definition, including a Private Enterprise Number (PEN) [IANA-PEN] to identify that the application registry being used is not owned by the Exporter manufacturer or to identify the original enterprise in the case of a mediator or 3rd party device. The Selector ID represents the enterprise unique global ID for layer 7 applications. The Selector ID has a global significance for all devices from the same enterprise.

Some values (5, 7, 8, 9, 10, 11, 14, 15, 16, and 17), are reserved to be compliant with existing implementations already using the classificationEngineId.

Abstract Data Type: unsigned8
Data Type Semantics: identifier
ElementId: 101
Status: current

7.1.5. applicationCategoryName

Name: applicationCategoryName
Description:
An attribute that provides a first-level categorization for each Application Id.
Abstract Data Type: string
Data Type Semantics:
ElementId: 372
Status: current

7.1.6. applicationSubCategoryName

Name: applicationSubCategoryName
Description:
An attribute that provides a second-level categorization for each Application Id.
Abstract Data Type: string
Data Type Semantics:
ElementId: 373
Status: current

7.1.7. applicationGroupName

Name: applicationGroupName
Description:
An attribute that groups multiple Application IDs that belong to the same networking application.
Abstract Data Type: string
Data Type Semantics:
ElementId: 374
Status: current

7.1.8. p2pTechnology

Name: p2pTechnology
Description:
Specifies if the Application ID is based on peer-to-peer technology. Possible values are { "yes", "y", 1 }, { "no", "n", 2 }, and { "unassigned", "u", 0 }.
Abstract Data Type: string
Data Type Semantics:
ElementId: 288
Status: current

7.1.9. tunnelTechnology

Name: tunnelTechnology
Description:
Specifies if the Application ID is used as a tunnel technology. Possible values are { "yes", "y", 1 }, { "no", "n", 2 }, and { "unassigned", "u", 0 }.
Abstract Data Type: string
Data Type Semantics:
ElementId: 289
Status: current

7.1.10. encryptedTechnology

Name: encryptedTechnology
Description:
Specifies if the Application ID is an encrypted networking protocol. Possible values are { "yes", "y", 1 }, { "no", "n", 2 }, and { "unassigned", "u", 0 }.
Abstract Data Type: string
Data Type Semantics:
ElementId: 290
Status: current

7.2. Classification Engine ID Registry

The Information Element #101, named classificationEngineId, carries information about the context for the Selector ID, and can be considered as a specific registry for application assignments. For ensuring extensibility of this information, IANA has created a new registry for Classification Engine IDs and filled it with the initial list from the description Information Element #101, classificationEngineId, along with their respective default lengths (Table 2 in this document).

New assignments for Classification Engine IDs will be administered by IANA through Expert Review [RFC5226], i.e., review by one of a group of experts designated by an IETF Area Director. The group of experts must double-check the new definitions with already defined Classification Engine IDs for completeness, accuracy, and redundancy. The specification of Classification Engine IDs MUST be published using a well-established and persistent publication medium.

8. Security Considerations

The same security considerations as for the IPFIX protocol [RFC5101] apply. The IPFIX extension specified in this memo allows to identify what applications are used on the network. Consequently, it is possible to identify what applications are being used by the users, potentially threatening the privacy of those users, if not handled with great care.

As mentioned in Section 1.1, the application knowledge is useful in security based applications. Security applications may impose supplementary requirements on the export of application information, and these need to be examined on a case by case basis.

9. References

9.1. Normative References

[ETHERTYPE] IEEE, <http://standards.ieee.org/develop/regauth/ethertype/eth.txt>.

[IANA-PEN] IANA, "PRIVATE ENTERPRISE NUMBERS", <http://www.iana.org/assignments/enterprise-numbers>.

[IANA-PORTS] IANA, "Service Name and Transport Protocol Port Number Registry", <http://www.iana.org/assignments/port-numbers>.

[IANA-PROTO] IANA, "Protocol Numbers", <http://www.iana.org/assignments/protocol-numbers>.

[LLC] IEEE, "LOGICAL LINK CONTROL (LLC) PUBLIC LISTING", <http://standards.ieee.org /develop/regauth/llc /public.html>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC5101] Claise, B., Ed., "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information", RFC 5101, January 2008.

[RFC5102] Quittek, J., Bryant, S., Claise, B., Aitken, P., and J. Meyer, "Information Model for IP Flow Information Export", RFC 5102, January 2008.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

9.2. Informative References

[CISCO-APPLICATION-REGISTRY]
Cisco, "Application Registry", <http://www.cisco.com/go/application_registry>.

[IANA-IPFIX] IANA, "IP Flow Information Export (IPFIX) Entities", <http://www.iana.org/assignments/ipfix>.

[LLDP] IEEE, Std 802.1AB-2005, "Standard for Local and metropolitan area networks - Station and Media Access Control Connectivity Discovery", IEEE Std 802.1AB-2005 IEEE Std, 2005.

[RFC792] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, September 1981.

[RFC3917] Quittek, J., Zseby, T., Claise, B., and S. Zander, "Requirements for IP Flow Information Export (IPFIX)", RFC 3917, October 2004.

[RFC3954] Claise, B., Ed., "Cisco Systems NetFlow Services Export Version 9", RFC 3954, October 2004.

[RFC4020] Kompella, K. and A. Zinin, "Early IANA Allocation of Standards Track Code Points", BCP 100, RFC 4020, February 2005.

[RFC5103] Trammell, B. and E. Boschi, "Bidirectional Flow Export Using IP Flow Information Export (IPFIX)", RFC 5103, January 2008.

[RFC5470] Sadasivan, G., Brownlee, N., Claise, B., and J. Quittek, "Architecture for IP Flow Information Export", RFC 5470, March 2009.

[RFC5471] Schmoll, C., Aitken, P., and B. Claise, "Guidelines for IP Flow Information Export (IPFIX) Testing", RFC 5471, March 2009.

[RFC5473] Boschi, E., Mark, L., and B. Claise, "Reducing Redundancy in IP Flow Information Export (IPFIX) and Packet Sampling (PSAMP) Reports", RFC 5473, March 2009.

[RFC5476] Claise, B., Ed., Johnson, A., and J. Quittek, "Packet Sampling (PSAMP) Protocol Specifications", RFC 5476, March 2009.

[RFC5477] Dietz, T., Claise, B., Aitken, P., Dressler, F., and G. Carle, "Information Model for Packet Sampling Exports", RFC 5477, March 2009.

[RFC5353] Xie, Q., Stewart, R., Stillman, M., Tuexen, M., and A. Silverton, "Endpoint Handlespace Redundancy Protocol (ENRP)", RFC 5353, September 2008.

[RFC5811] Hadi Salim, J. and K. Ogawa, "SCTP-Based Transport Mapping Layer (TML) for the Forwarding and Control Element Separation (ForCES) Protocol", RFC 5811, March 2010.

[RFC6183] Kobayashi, A., Claise, B., Muenz, G., and K. Ishibashi, "IP Flow Information Export (IPFIX) Mediation: Framework", RFC 6183, April 2011.

[RFC6313] Claise, B., Dhandapani, G., Aitken, P., and S. Yates, "Export of Structured Data in IP Flow Information Export (IPFIX)", RFC 6313, July 2011.

10. Acknowledgements

The authors would like to thank their many colleagues across Cisco Systems who made this work possible. Specifically, Patrick Wildi for his time and expertise.

Appendix A. Additions to XML Specification of IPFIX Information Elements (Non-normative)

This appendix contains additions to the machine-readable description of the IPFIX information model coded in XML in Appendix A and Appendix B in [RFC5102]. Note that this appendix is of informational nature, while the text in Section 7 (generated from this appendix) is normative.

The following field definitions are appended to the IPFIX information model in Appendix A of [RFC5102].


```
  <field name="applicationDescription"
         dataType="string"
         group="application"
         elementId="94" applicability="all"
status="current">
    <description>
      <paragraph>
         Specifies the description of an application.
      </paragraph>
    </description>
  </field>

  <field name="applicationId"
         dataType="octetArray"
         group="application"
         dataTypeSemantics="identifier"
         elementId="95" applicability="all"
status="current">
    <description>
      <paragraph>
         Specifies an Application ID.
      </paragraph>
    </description>
    <reference>
      <paragraph>
         See Section 4 of [RFC6759]
        for the applicationId Information Element
        Specification.
      </paragraph>
    </reference>
  </field>

  <field name="applicationName"
         dataType="string"
         group="application"
         elementId="96" applicability="all"
```

```
   status="current">
       <description>
         <paragraph>
            Specifies the name of an application.
         </paragraph>
       </description>
     </field>

     <field name="classificationEngineId"
            dataType="unsigned8"
            group="application"
            dataTypeSemantics="identifier"
            elementId="101" applicability="all"
   status="current">
       <description>
         <paragraph>
           0 Invalid.

           1 IANA-L3: The Assigned Internet Protocol Number
             (layer 3 (L3)) is exported in the Selector ID.
             See http://www.iana.org/assignments/protocol-
             numbers.

           2 PANA-L3: Proprietary layer 3 definition.  An
             enterprise can export its own layer 3 protocol
             numbers.  The Selector ID has a global
             significance for all devices from the same
             enterprise.

           3 IANA-L4: The IANA layer 4 (L4) well-known port
             number is exported in the Selector ID.  See
             [IANA-PORTS].  Note: as an IPFIX flow is
             unidirectional, it contains the destination
             port.

           4 PANA-L4: Proprietary layer 4 definition.  An
             enterprise can export its own layer 4 port
             numbers.  The Selector ID has global
             significance for devices from the same
             enterprise.  Example: IPFIX was pre-assigned
             port 4739 using the IANA early allocation
             process [RFC4020] years before the document was
             published as an RFC.  While waiting for the
             RFC and its associated IANA registration,
             Selector ID 4739 was used with this PANA-L4.

           5 Reserved
```

6 USER-Defined: The Selector ID represents applications defined by the user (using CLI, GUI, etc.) based on the methods described in Section 2. The Selector ID has a local significance per device.

7 Reserved

8 Reserved

9 Reserved

10 Reserved

11 Reserved

12 PANA-L2: Proprietary layer 2 (L2) definition. An enterprise can export its own layer 2 identifiers. The Selector ID represents the enterprise's unique global layer 2 applications. The Selector ID has a global significance for all devices from the same enterprise. Examples include the Cisco Subnetwork Access Protocol (SNAP).

13 PANA-L7: Proprietary layer 7 definition. The Selector ID represents the enterprise's unique global ID for layer 7 applications. The Selector ID has a global significance for all devices from the same enterprise. This Classification Engine ID is used when the application registry is owned by the Exporter manufacturer (referred to as the "enterprise" in this document).

14 Reserved

15 Reserved

16 Reserved

17 Reserved

18 ETHERTYPE: The Selector ID represents the well-known Ethertype. See [ETHERTYPE].

19 LLC: The Selector ID represents the well-known IEEE 802.2 Link Layer Control (LLC)

```
             Destination Service Access Point (DSAP).  See
             [LLC].

          20 PANA-L7-PEN: Proprietary layer 7 definition,
             including a Private Enterprise Number (PEN)
             [IANA-PEN] to identify that the application
             registry being used is not owned by the
             Exporter manufacturer or to identify the
             original enterprise in the case of a mediator
             or 3rd party device.  The Selector ID represents
             the enterprise unique global ID for layer 7
             applications.  The Selector ID has a global
             significance for all devices from the same
             enterprise.
        </paragraph>
      </description>
    </field>

    <field name="applicationCategoryName"
           dataType="string"
           group="application"
           elementId="372"
           applicability="all"
           status="current">
      <description>
        <paragraph>
           An attribute that provides a first-level categorization
           for each Application Id.
        </paragraph>
      </description>
    </field>

    <field name="applicationSubCategoryName"
           dataType="string"
           group="application"
           elementId="373"
           applicability="all"
           status="current">
      <description>
        <paragraph>
           An attribute that provides a second-level
           categorization for each Application ID.
        </paragraph>
      </description>
    </field>

    <field name="applicationGroupName"
           dataType="string"
```

           group="application"
           elementId="374"
           applicability="all"
           status="current">
      <description>
        <paragraph>
           An attribute that groups multiple Application IDs
           that belong to the same networking application.
        </paragraph>
      </description>
    </field>

    <field name="p2pTechnology"
           dataType="string"
           group="application"
           elementId="288"
           applicability="all"
           status="current">
      <description>
        <paragraph>
           Specifies if the Application ID is based on peer-
           to-peer technology. Possible values are
           { "yes", "y", 1 }, { "no", "n", 2 }, and
           { "unassigned", "u", 0 }.
        </paragraph>
      </description>
    </field>

    <field name="tunnelTechnology"
           dataType="string"
           group="application"
           elementId="289"
           applicability="all"
           status="current">
      <description>
        <paragraph>
           Specifies if the Application ID is used as a
           tunnel technology. Possible values are
           { "yes", "y", 1 }, { "no", "n", 2 }, and
           { "unassigned", "u", 0 }.
        </paragraph>
      </description>
    </field>

    <field name="encryptedTechnology"
           dataType="string"
           group="application"
           elementId="290"

```
            applicability="all"
            status="current">
       <description>
         <paragraph>
            Specifies if the Application ID is an encrypted
            networking protocol.  Possible values are
            { "yes", "y", 1 }, { "no", "n", 2 }, and
            { "unassigned", "u", 0 }.
         </paragraph>
       </description>
     </field>
```


Appendix B. Port Collisions Tables (Non-normative)

The following table lists the 10 ports that have different protocols assigned for TCP and UDP (at the time of writing this document):

exec	512/tcp	remote process execution; authentication performed using passwords and UNIX login names
comsat/biff	512/udp	used by mail system to notify users of new mail received; currently receives messages only from processes on the same machine
login	513/tcp	remote login a la telnet; automatic authentication performed based on priviledged [sic] port numbers and distributed data bases which identify "authentication domains"
who	513/udp	maintains data bases showing who’s logged in to machines on a local net and the load average of the machine
shell	514/tcp	cmd like exec, but automatic authentication is performed as for login server

syslog	514/udp	
oob-ws-https	664/tcp	DMTF out-of-band secure web services management protocol Jim Davis <jim.davis@wbemsolutions.com>
asf-secure-rmcp	664/udp	ASF Secure Remote Management and Control Protocol
rfile	750/tcp	
kerberos-iv	750/udp	kerberos version iv
submit	773/tcp	
notify	773/udp	
rpasswd	774/tcp	
acmaint_dbd	774/udp	
entomb	775/tcp	
acmaint_transd	775/udp	
busboy	998/tcp	
puparp	998/udp	
garcon	999/tcp	
applix	999/udp	Applix ac

Table 4: Different Protocols on UDP and TCP

The following table lists the 19 ports that have different protocols assigned for TCP and SCTP (at the time of writing this document):

#	3097/tcp	Reserved
itu-bicc-stc	3097/sctp	ITU-T Q.1902.1/Q.2150.3 Greg Sidebottom <gregside@home.com>
#	5090/tcp	<not assigned>
car	5090/sctp	Candidate AR
#	5091/tcp	<not assigned>
cxtp	5091/sctp	Context Transfer Protocol

```
      #                6704/tcp    Reserved

      frc-hp           6704/sctp   ForCES HP (High Priority)
                                   channel [RFC5811]

      #                6705/tcp    Reserved

      frc-mp           6705/sctp   ForCES MP (Medium
                                   Priority) channel
                                   [RFC5811]

      #                6706/tcp    Reserved

      frc-lp           6706/sctp   ForCES LP (Low Priority)
                                   channel [RFC5811]

      #                9082/tcp    <not assigned>

      lcs-ap           9082/sctp   LCS Application Protocol
                                   Kimmo Kymalainen
                                   <kimmo.kymalainen@etsi.org>

      #                9902/tcp    <not assigned>

      enrp-sctp-tls    9902/sctp   enrp/tls server channel
                                   [RFC5353]

      #                11997/tcp   <not assigned>
      #                11998/tcp   <not assigned>
      #                11999/tcp   <not assigned>

      wmereceiving     11997/sctp  WorldMailExpress
      wmedistribution  11998/sctp  WorldMailExpress
      wmereporting     11999/sctp  WorldMailExpress
                                Greg Foutz
                                   <gregf@adminovation.com>

      #                25471/tcp   <not assigned>

      rna              25471/sctp  RNSAP User Adaptation for
                                   Iurh
                                   Dario S. Tonesi
                                   <dario.tonesi@nsn.com>
                                   07 February 2011

      #                29118/tcp   Reserved

      sgsap            29118/sctp  SGsAP in 3GPP
```

#	29168/tcp	Reserved
sbcap	29168/sctp	SBcAP in 3GPP
#	29169/tcp	<not assigned>
iuhsctpassoc	29169/sctp	HNBAP and RUA Common Association John Meredith <John.Meredith@etsi.org> 08 September 2009
#	36412/tcp	<not assigned>
s1-control	36412/sctp	S1-Control Plane (3GPP) Kimmo Kymalainen <kimmo.kymalainen@etsi.org> 01 September 2009
#	36422/tcp	<not assigned>
x2-control	36422/sctp	X2-Control Plane (3GPP) Kimmo Kymalainen <kimmo.kymalainen@etsi.org> 01 September 2009
#	36443/tcp	<not assigned>
m2ap	36443/sctp	M2 Application Part Dario S. Tonesi <dario.tonesi@nsn.com> 07 February 2011
#	36444/tcp	<not assigned>
m3ap	36444/sctp	M3 Application Part Dario S. Tonesi <dario.tonesi@nsn.com> 07 February 2011

Table 5: Different Protocols on SCTP and TCP

Appendix C. Application Registry Example (Non-normative)

A reference to the Cisco Systems assigned numbers for the Application ID and the different attribute assignments can be found at [CISCO-APPLICATION-REGISTRY].

Authors' Addresses

Benoit Claise
Cisco Systems, Inc.
De Kleetlaan 6a b1
Diegem 1813
Belgium

Phone: +32 2 704 5622
EMail: bclaise@cisco.com

Paul Aitken
Cisco Systems, Inc.
96 Commercial Quay
Commercial Street
Edinburgh, EH6 6LX
United Kingdom

Phone: +44 131 561 3616
EMail: paitken@cisco.com

Nir Ben-Dvora
Cisco Systems, Inc.
32 HaMelacha St.,
P.O. Box 8735, I.Z.Sapir
South Netanya, 42504
Israel

Phone: +972 9 892 7187
EMail: nirbd@cisco.com