

IPFIX Working Group
Internet-Draft
Intended Status: Informational
Expires: August 20, 2012

B. Claise
P. Aitken
N. Ben-Dvora
Cisco Systems, Inc.
February 20, 2012

Export of Application Information in IPFIX
draft-claise-export-application-info-in-ipfix-05

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

This Internet-Draft will expire on August 20, 2012.

Internet-Draft <Export of App. Info. in IPFIX > Feb 2012

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in [Section 4.e](#) of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Abstract

This document specifies an extension to the IPFIX information model specified in [[RFC5102](#)] to export application information.

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](#)].

Internet-Draft <Export of App. Info. in IPFIX >

Feb 2012

Table of Contents

1.	Overview.....	4
1.1.	IPFIX Documents Overview.....	4
2.	Introduction.....	5
2.1.	Application Information Use Cases.....	7
3.	Terminology.....	8
3.1.	New Terminology.....	8
4.	applicationId Information Element Specification.....	8
4.1.	Existing Classification Engine IDs.....	9
4.2.	Selector ID Length per Classification IDs.....	12
4.3.	Application Name Options Template Record.....	13
4.4.	Resolving IANA L4 port collisions.....	14
5.	Grouping the Applications with the Attributes.....	19
5.1.	Options Template Record for the Attribute Values... 	21
6.	Application Id Examples.....	21
6.1.	Example 1: Layer 2 Protocol.....	21
6.2.	Example 2: Standardized IANA Layer 3 Protocol.....	22
6.3.	Example 3: Proprietary Layer 3 Protocol.....	23
6.4.	Example 4: Standardized IANA Layer 4 Port.....	25
6.5.	Example 4: Layer 7 Application.....	26
6.6.	Example: port Obfuscation.....	28
6.7.	Example: Application Mapping Options Template.....	29
6.8.	Example: Attributes Values Options Template Record.	30
7.	IANA Considerations.....	31
7.1.	New Information Elements.....	31
7.1.1.	applicationDescription.....	31
7.1.2.	applicationId.....	31
7.1.3.	applicationName.....	32
7.1.4.	classificationEngineId.....	32
7.1.5.	applicationCategoryName.....	34
7.1.6.	applicationSubCategoryName.....	34
7.1.7.	applicationGroupName.....	35

7.1.8. p2pTechnology.....	35
7.1.9. tunnelTechnology.....	35
7.1.10. encryptedTechnology.....	35
7.2. Classification Engine Ids Registry.....	36
8. Security Considerations.....	36
9. References.....	36
9.1. Normative References.....	36
9.2. Informative References.....	37
10. Acknowledgement.....	38
11. Authors' Addresses.....	39
Appendix A. Additions to XML Specification of IPFIX	
Information Elements.....	39

List of Figures and Tables

Figure 1: applicationId Information Element	8
Table 1: Existing Classification Engine IDs	12
Table 2: Selector ID default length per Classification Engine ID	13
Table 3: IANA layer 4 port collisions between UDP and TCP .	16
Table 4: IANA layer 4 port collisions between SCTP and TCP	19
Table 5: Existing Application Id Static Attributes	20

[1. Overview](#)

[1.1. 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 name, type 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 IP Flow Information Export (IPFIX) protocol, representing each Biflow using a single Flow Record.

The "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.

[2](#). 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 information elements, 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 protocols (such as ARP, PPP, LLDP)
2. IP protocols (such as ICMP, IGMP, GRE)
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 with a minimum 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 different port than the IANA assigned one. For example an HTTP server might run a TCP port 23 (assigned to telnet in [[IANA-PORTS](#)])

2. IANA does 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. Detecting applications based only on port matching might wrongly identify the traffic. Note that this example stresses the need for the engine strength. 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 in use must be exported using the destination port (destinationTransportPort at [IANA-IPFIX]) in the corresponding IPFIX record.

This document specifies the Application Id (as described in [section 4](#).) to export the application information with the IPFIX protocol [[RFC5101](#)].

Application could be defined at different OSI layers, from layer 2 to layer 7. For example: Link Layer Distribution Protocol (LLDP) [[LLDP](#)] is layer 2 application, ICMP is layer 3 application [[IANA-PROTO](#)], HTTP is layer 4 application [[IANA-PORTS](#)], and skype is 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, for example Skype, are not available. Some reverse engineering as well as a ubiquitous language for application identification, would be two required conditions

to be able to manage an IANA registry for these types of applications. Clearly, these are blocking factors. As this specification focuses on the application information encoding, this document doesn't contain an application registry for non IANA applications. However, a reference to the Cisco Systems assigned numbers for the Application Id and the different attribute assignments can be found at [[CISCO](#)]. Note that, while the Application Ids for all applications are available at [[CISCO](#)], only the layer 7 ones cannot be compiled from existing well known registries.

[2.1.](#) Application Information Use Cases

There are several use cases for application information:

1. Application Visibility

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

2. Billing Services

In some cases, network providers are willing to bill different applications differently. For example, provide different billing for VoIP and Web browsing.

3. Congestion Control

While traffic demand is increasing (mainly due to the high usage of peer to peer applications, video applications and web download applications), the providers revenue doesn't grow. Providers are looking at a more efficient way to control and prioritize the network utilization. An application aware bandwidth control system is used to prioritize the traffic based on the applications, giving the critical applications priority over the non-critical applications.

4. 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.

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

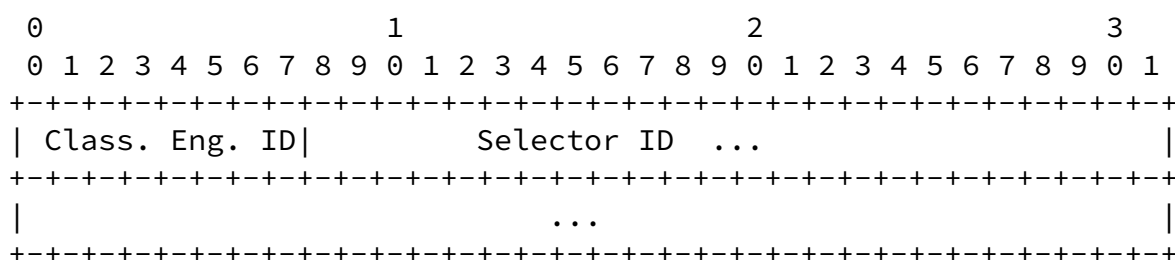


Figure 1: applicationId Information Element

Classification Engine ID

A unique identifier for the engine which 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 in sync with the selectorId Information Element, specified in the PSAMP Protocol [[RFC5476](#)].

[4.1](#). Existing Classification Engine IDs

The following Classification Engine IDs have been allocated:

Name	Value	Description
	0	Invalid.
IANA-L3	1	The IANA protocol (layer 3) number is exported in the Selector ID. See [IANA-PROTO].
PANA-L3	2	Proprietary layer 3 definition. A company can export its own layer 3 protocol numbers, while waiting for IANA to assign it. The Selector ID has a global significance for all devices from the same company. Hopefully the same Selector IDs will be maintained after the IANA standardization.
IANA-L4	3	The IANA layer 4 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 in a flow from the client to the server.

PANA-L4	4	Proprietary layer 4 definition. A company can export its own layer 4 port numbers, while waiting for IANA to assign it. The Selector ID has global significance for devices from the same company. Hopefully the same Selector IDs will be maintained after the IANA standardization. Example: IPFIX had the port 4739 pre-assigned in the IETF draft for years. While waiting for the RFC and its associated IANA registration, the Selector ID 4739 was used with this PANA-L4.
	5	Reserved.
	6	The Selector ID represents applications defined by the user (using CLI or GUI) based on the methods described in section 2 . The Selector ID has a local significance per device.
USER-Defined	7	Reserved.
	8	Reserved.
	9	Reserved.
	10	Reserved.
	11	Reserved.
PANA-L2	12	Proprietary layer 2 definition. A company can export its own layer 2 identifiers. The Selector ID represents the company unique global layer 2 applications. The Selector ID has

a global significance for all devices from the same company.

Internet-Draft <Export of App. Info. in IPFIX >

Feb 2012

Examples include Cisco Subnetwork Access Protocol (SNAP).

PANA-L7	13	Proprietary layer 7 definition. The Selector ID represents the company unique global ID for the layer 7 applications. The Selector ID has a global significance for all devices from the same company. A reference to the Cisco Systems assigned numbers for the layer 7 Application Id assignments can be found at [CISCO].
	14	Reserved.
	15	Reserved.
	16	Reserved.
	17	Reserved.
ETHERTYPE	18	The Selector ID represents the well-known Ethertype. See [ETHERTYPE]. Note that the Ethertype is usually expressed in hexadecimal. However, the corresponding decimal value is used in this Selector ID.
LLC	19	The Selector ID represents the well-known IEEE 802.2 Link Layer Control (LLC) Destination Service Access Point (DSAP). See [LLC]. Note that LLC DSAP is usually expressed in hexadecimal. However, the corresponding decimal value is used in this

Selector ID.

20 to
254 Available.

MAX 255 255 is the maximum Engine ID.

<Claise, Aitken, Ben-Dvora>

Expires Aug 20 2012

[Page 11]

Internet-Draft <Export of App. Info. in IPFIX > Feb 2012

Table 1: Existing Classification Engine IDs

Note 1: "PANA = Proprietary Assigned Number Authority". In other words, a company specific version of IANA for internal IDs.

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 these specifications. However, this Information Element was created for completeness.

[4.2](#). Selector ID Length per Classification IDs

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 the IPFIX export.

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

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 bytes for Cisco Subnetwork Access Protocol (SNAP).
PANA-L7	3
ETHERTYPE	2

LLC	1
-----	---

Table 2: Selector ID default length
per Classification Engine 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 the 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 protocol field from the IP header encoded in one byte. However, an IANA L3 protocol encoding may be encoded with 3 bytes. In such a case, the Selector ID value MUST always be encoded in the least significant bits as shown in Figure 2.




Figure 2: Selector ID encoding

[4.3.](#) Application Name Options Template Record

For Classification Engines which 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 which 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).

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. The mechanism for the Collector to know about Exporter enterprise IDs is out of scope of this document. Possible tracks are: SNMP polling, an Options Template export, hardcoded value, etc.

[4.4.](#) Resolving IANA L4 port collisions

Even though the IANA L4 ports usually point to the same protocols for both UDP, TCP or other transport types, there are some exceptions. 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 privileged port numbers and distributed data
bases		which identify "authentication domains"

<Claire, Aitken, Ben-Dvora>

Expires Aug 20 2012

[Page 14]

Internet-Draft <Export of App. Info. in IPFIX >

Feb 2012

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 web	664/tcp	DMTF out-of-band secure services management protocol Jim Davis
<jim.davis@wbemsolutions.com>		June 2007
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

<Claire, Aitken, Ben-Dvora>

Expires Aug 20 2012

[Page 15]

Internet-Draft <Export of App. Info. in IPFIX >

Feb 2012

Table 3: IANA layer 4 port collisions between 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>
cxtcp Protocol	5091/sctp	Context Transfer RFC 4065 - July 2005
#	6704/tcp	Reserved
frc-hp Priority)	6704/sctp	ForCES HP (High 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>
		04 June 2010
#	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> March 2006
#	25471/tcp	<not assigned>
rna for	25471/sctp	RNSAP User Adaptation
		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>
ihusctpassoc	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) KimmoKymalainen
<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 4: IANA layer 4 port collisions between SCTP and TCP

Instead of imposing the transport protocol (UDP/TCP/SCTP/etc.) in the scope of the "Application Name Options Template Record" for all applications (on top of having the transport protocol as 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 DCCP.

[5. Grouping the Applications with the 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 contain the 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 5: Existing Application Id Static Attributes

Every application is assigned to one ApplicationCategoryName, one ApplicationSubCategoryName, one ApplicationGroupName, has one p2pTechnology, one tunnelTechnology, and one encryptedTechnology.

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

company specific. The Cisco Systems attribute values for the different applications are available at [[CISCO](#)].

[5.1](#). Options Template Record for the 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 ApplicationId 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 6.8](#). below.

[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, 18, and 19.

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 while the Selector ID has the value 35020.

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.

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 is 1:

IANA-	1	The IANA protocol (layer 3) number is exported in the Selector ID.
L3		See [IANA-PROTO].

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

Decimal	Keyword	Protocol	Reference
---------	---------	----------	-----------

Internet-Draft <Export of App. Info. in IPFIX >

Feb 2012

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 know registry, ie the IANA L3 protocol number.

[6.3](#). Example 3: Proprietary Layer 3 Protocol

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

<Claise, Aitken, Ben-Dvora>

Expires Aug 20 2012

[Page 23]

Internet-Draft <Export of App. Info. in IPFIX >

Feb 2012

From the list of Classification Engine IDs in Table 1, the proprietary layer 3 Classification Engine ID is 2:

PANA- L3	2	Proprietary layer 3 definition. A company can export its own layer 3 protocol numbers, while waiting for IANA to assign it. The Selector ID has a global significance for all devices from the same company. Hopefully the same Selector IDs will be maintained after the IANA standardization.
-------------	---	---

A global registry within the company 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 company, 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)

<Claise, Aitken, Ben-Dvora>

Expires Aug 20 2012

[Page 24]

Internet-Draft <Export of App. Info. in IPFIX >

Feb 2012

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 6.7](#).

[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 is 3:

IANA- L4	3	The IANA layer 4 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 in a flow from the client to the server.
-------------	---	---

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 '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)
- 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 know registry, ie the IANA L4 protocol number.

6.5. Example 4: Layer 7 Application

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

From the list of Classification Engine IDs in Table 1, the L7 unique Classification Engine ID is 13:

PANA-L7	13	Proprietary layer 7 definition. The Selector ID represents the company unique global ID for the layer 7 applications. The Selector ID has a global significance for all devices from the same company. A reference to the Cisco Systems assigned numbers for the layer 7 Application Id assignments can be found at [CISCO] .
---------	----	---

[Page 26]

Internet-Draft <Export of App. Info. in IPFIX > Feb 2012

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

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

Therefore the Application Id is encoded as:

[illegible]

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 company, so that the Collector can determine which application is represented by the Application Id by loading the registry out of band. A reference to the Cisco Systems assigned numbers for the layer 7 Application Id and the different attribute assignments can be found at [[CISCO](#)].

Along with this Flow Record, a new Options Template Record would be exported, as shown in [Section 6.7](#).

[6.6](#). 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 the default HTTP port 80 or 8080 was used, it must export the destination port (destinationTransportPort at [[IANA-IPFIX](#)]) in the corresponding IPFIX record.

In the case of a standardized IANA layer 4 port, the Classification Engine ID is 2, and the Selector ID has the

value of 80 for HTTP (see [[IANA-PORTS](#)]).

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.

[6.7](#). Example: Application Mapping Options Template

Along with the Flow Records shown in the above examples, a new Options Template Record would 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="Citrix",  
  applicationDescription="A Citrix 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 "Citrix" application.

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: the applicationDescription, applicationId, applicationName, classificationEngineId, applicationCategoryName, applicationSubCategoryName, applicationGroupName, p2pTechnology, tunnelTechnology, and encryptedTechnology.

New Information Elements to be added to the IPFIX Information Element registry at [IANA-IPFIX] are listed below.

EDITOR'S NOTE: the XML specification in [Appendix A](#) must be updated with the elementID values allocated below.

[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 [EDITORS NOTE: this document] 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 which 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.

0 Invalid.

- 1 IANA-L3: The IANA protocol (layer 3) number is exported in the Selector ID. See <http://www.iana.org/assignments/protocol-numbers>.
- 2 PANA-L3: Proprietary layer 3 definition. A company can export its own layer 3 protocol numbers, while waiting for IANA to assign it. The Selector ID has a global significance for all devices from the same company. Hopefully the same Selector IDs will be maintained after the IANA standardization.
- 3 IANA-L4: The IANA layer 4 well-known port number is exported in the Selector ID. See <http://www.iana.org/assignments/port-numbers>. Note: as an IPFIX flow is unidirectional, it contains the destination port in a flow from the client to the server.

Internet-Draft <Export of App. Info. in IPFIX >

Feb 2012

- 4 PANA-L4: Proprietary layer 4 definition. A company can export its own layer 4 port numbers, while waiting for IANA to assign it. The Selector ID has global significance for devices from the same company. Hopefully the same Selector IDs will be maintained after the IANA standardization. Example: IPFIX had the port 4739 pre-assigned in the IETF draft for years. While waiting for the RFC and its associated IANA registration, the 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 or GUI) 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 definition. A company can export its own layer 2 identifiers. The Selector ID represents the company unique global layer 2 applications. The Selector ID has a global significance for all devices from the same company. Examples include Cisco Subnetwork Access Protocol (SNAP).
- 13 PANA-L7: Proprietary layer 7 definition. The Selector ID represents the company unique global ID for the layer 7 applications. The Selector ID has a global significance for all devices from the same

company.

14 Reserved

15 Reserved

16 Reserved

<Claise, Aitken, Ben-Dvora>

Expires Aug 20 2012

[Page 33]

Internet-Draft <Export of App. Info. in IPFIX >

Feb 2012

17 Reserved

18 ETHERTYPE: The Selector ID represents the well-known Ethertype. See <http://standards.ieee.org/develop/regauth/ethertype/eth.txt>. Note that the Ethertype is usually expressed in hexadecimal. However, the corresponding decimal value is used in this Selector ID.

19 LLC: The Selector ID represents the well-known IEEE 802.2 Link Layer Control (LLC) Destination Service Access Point (DSAP). See <http://standards.ieee.org/develop/regauth/ethertype/eth.txt>. Note that LLC DSAP is usually expressed in hexadecimal. However, the corresponding decimal value is used in this Selector ID.

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: <to be assigned>

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: <to be assigned>

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: <to be assigned>

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 }.

<Claise, Aitken, Ben-Dvora>

Expires Aug 20 2012

[Page 35]

Internet-Draft <Export of App. Info. in IPFIX >

Feb 2012

Abstract Data Type: string

Data Type Semantics:

ElementId: 290

Status: current

[7.2](#). Classification Engine Ids 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.

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.

RFC-EDITOR: this should be assigned similarly to mplsTopLabelType subregistry at <http://www.iana.org/assignments/ipfix/ipfix.xml>

8. Security Considerations

The same security considerations as for the IPFIX Protocol [[RFC5101](#)] apply.

9. References

9.1. Normative References

[RFC2119] S. Bradner, 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.

<Claise, Aitken, Ben-Dvora>

Expires Aug 20 2012

[Page 36]

Internet-Draft <Export of App. Info. in IPFIX > Feb 2012

[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. Alverstrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [RFC 5226](#), May 2008

9.2. Informative References

[RFC792] J. Postel, Internet Control Message Protocol, [RFC 792](#), September 1981.

[RFC3917] Quittek, J., Zseby, T., Claise, B., and S. Zander,

Requirements for IP Flow Information Export, [RFC 3917](#), October 2004.

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

[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., "Packet Sampling (PSAMP) Protocol Specifications", [RFC 5476](#), March 2009.

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

<Claise, Aitken, Ben-Dvora>

Expires Aug 20 2012

[Page 37]

Internet-Draft <Export of App. Info. in IPFIX >

Feb 2012

[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.

[IANA-IPFIX] <http://www.iana.org/assignments/ipfix/ipfix.xml>

[IANA-PORTS] <http://www.iana.org/assignments/port-numbers>

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

[ETHERTYPE]

<http://standards.ieee.org/develop/regauth/ethertype/eth.txt>

[LLC]

<http://standards.ieee.org/develop/regauth/llc/public.html>.

[CISCO] <http://www.cisco.com>

10. Acknowledgement

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.

11. 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

[Appendix A](#). Additions to XML Specification of IPFIX Information Elements

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 [EDITORS NOTE: this document] 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">

```

Internet-Draft <Export of App. Info. in IPFIX >

Feb 2012

<description>

<paragraph>

A unique identifier for the engine which 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.

0 Invalid.

1 IANA-L3: The IANA protocol (layer 3) number is exported in the Selector ID. See <http://www.iana.org/assignments/protocol-numbers>.

2 PANA-L3: Proprietary layer 3 definition. A company can export its own layer 3 protocol numbers, while waiting for IANA to assign it. The Selector ID has a global significance for all devices from the same company. Hopefully the same Selector IDs will be maintained after the IANA standardization.

3 IANA-L4: The IANA layer 4 well-known port number is exported in the Selector ID. See <http://www.iana.org/assignments/port-numbers>. Note: as an IPFIX flow is unidirectional, it contains the destination port in a flow from the client to the server.

4 PANA-L4: Proprietary layer 4 definition. A company can export its own layer 4 port numbers, while waiting for IANA to assign it. The Selector ID has global significance for devices from the same company. Hopefully the same Selector IDs will be maintained after the IANA standardization. Example: IPFIX had the port 4739 pre-assigned in the IETF draft for years. While

waiting for the RFC and its associated IANA registration, the 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 or GUI) 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 definition. A company can export its own layer 2 identifiers. The Selector ID represents the company unique global layer 2 applications. The Selector ID has a global significance for all devices from the same company. Examples include Cisco Subnetwork Access Protocol (SNAP).

13 PANA-L7: Proprietary layer 7 definition. The Selector ID represents the company unique global ID for the layer 7 applications. The Selector ID has a global significance for all devices from the same company.

14 Reserved

15 Reserved

16 Reserved

17 Reserved

18 ETHERTYPE: The Selector ID represents the well-known Ethertype. See <http://standards.ieee.org/develop/regauth/ethertype/eth.txt>. Note that the Ethertype is usually expressed in hexadecimal. However, the corresponding decimal value is used in this Selector ID.

<Claise, Aitken, Ben-Dvora>

Expires Aug 20 2012

[Page 42]

Internet-Draft <Export of App. Info. in IPFIX >

Feb 2012

19 LLC: The Selector ID represents the well-known IEEE 802.2 Link Layer Control (LLC) Destination Service Access Point (DSAP). See <http://standards.ieee.org/develop/regauth/ethertype/eth.txt>. Note that LLC DSAP is usually expressed in hexadecimal. However, the corresponding decimal value is used in this Selector ID.

</paragraph>

</description>

</field>

<field name="applicationCategoryName"

 dataType="string"

 group="application"

 elementId="<to be assigned>"

 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="<to be assigned>"
        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="<to be assigned>"
  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>

```

<Claise, Aitken, Ben-Dvora>

Expires Aug 20 2012

[Page 45]