

Network Working Group
Internet-Draft
Intended status: Informational
Expires: December 28, 2014

M. Tuexen, Ed.
Muenster Univ. of Appl. Sciences
F. Risso
Politecnico di Torino
J. Bongertz
Airbus DS CyberSecurity
G. Harris

June 26, 2014

PCAP Next Generation (PCAPNG) Dump File Format
draft-tuexen-opswg-pcapng-00.txt

Abstract

This document describes a format to dump captured packets to a file. This format is extensible; Wireshark can currently read and write it, and libpcap can currently read some pcap-ng files.

Status of This Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on December 28, 2014.

Copyright Notice

Copyright (c) 2014 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.

Table of Contents

1.	Introduction	3
2.	Conventions	3
3.	General File Structure	3
3.1.	General Block Structure	3
3.2.	Block Types	4
3.3.	Logical Block Hierarchy	6
3.4.	Physical File Layout	6
3.5.	Options	7
3.6.	Data format	10
4.	Block Definition	11
4.1.	Section Header Block (mandatory)	11
4.2.	Interface Description Block (mandatory)	14
4.3.	Enhanced Packet Block (optional)	23
4.4.	Simple Packet Block (optional)	27
4.5.	Packet Block (obsolete!)	29
4.6.	Name Resolution Block (optional)	31
4.7.	Interface Statistics Block (optional)	34
5.	Experimental Blocks (deserved to a further investigation) . .	37
5.1.	Alternative Packet Blocks (experimental)	37
5.2.	Compression Block (experimental)	37
5.3.	Encryption Block (experimental)	38
5.4.	Fixed Length Block (experimental)	39
5.5.	Directory Block (experimental)	40
5.6.	Traffic Statistics and Monitoring Blocks (experimental) .	41
5.7.	Event/Security Block (experimental)	41
6.	Recommended File Name Extension: .pcapng	41
7.	How to add Vendor / Domain specific extensions	41
8.	Conclusions	42
9.	Security Considerations	42
10.	IANA Considerations	42
11.	Acknowledgments	42
12.	References	42
12.1.	Normative References	42
12.2.	URIs	42
Appendix A.	Packet Block Flags Word	43
Appendix B.	Standardized Block Type Codes	43
	Authors' Addresses	44

1. Introduction

The problem of exchanging packet traces becomes more and more critical every day; unfortunately, no standard solutions exist for this task right now. One of the most accepted packet interchange formats is the one defined by libpcap, which is rather old and is lacking in functionality for more modern applications particularly from the extensibility point of view.

This document proposes a new format for dumping packet traces. The following goals are being pursued:

Extensibility: It should be possible to add new standard capabilities to the file format over time, and third parties should be able to enrich the information embedded in the file with proprietary extensions, with tools unaware of newer extensions being able to ignore them.

Portability: A capture trace must contain all the information needed to read data independently from network, hardware and operating system of the machine that made the capture.

Merge/Append data: It should be possible to add data at the end of a given file, and the resulting file must still be readable.

2. Conventions

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

3. General File Structure

3.1. General Block Structure

A capture file is organized in blocks, that are appended one to another to form the file. All the blocks share a common format, which is shown in Figure 1.

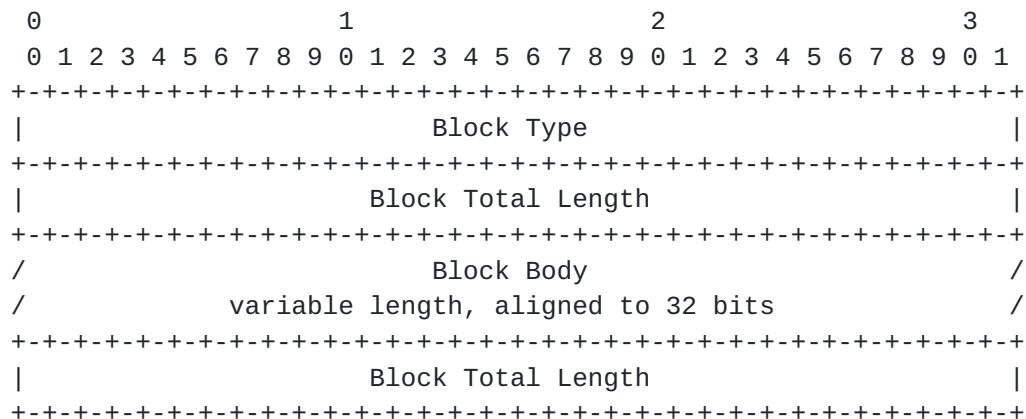


Figure 1: Basic block structure.

The fields have the following meaning:

- o Block Type (32 bits): unique value that identifies the block. Values whose Most Significant Bit (MSB) is equal to 1 are reserved for local use. They can be used to make extensions to the file format to save private data to the file. The list of currently defined types can be found in [Appendix B](#)
- o Block Total Length: total size of this block, in bytes. For instance, the length of a block that does not have body is 12 bytes.
- o Block Body: content of the block.
- o Block Total Length: total size of this block, in bytes. This field is duplicated for permitting backward file navigation.

This structure, shared among all blocks, makes it easy to process a file and to skip unneeded or unknown blocks. Some blocks can contain other blocks inside (nested blocks). Some of the blocks are mandatory, i.e. a dump file is not valid if they are not present, other are optional.

The General Block Structure allows defining other blocks if needed. A parser that does not understand them can simply ignore their content.

3.2. Block Types

The currently standardized Block Type codes are specified in [Appendix B](#), they have been grouped in the following four categories:

MANDATORY blocks MUST appear at least once in each file:

- o Section Header Block ([Section 4.1](#)): it defines the most important characteristics of the capture file.
- o Interface Description Block ([Section 4.2](#)): it defines the most important characteristics of the interface(s) used for capturing traffic.

OPTIONAL blocks MAY appear in a file:

- o Enhanced Packet Block ([Section 4.3](#)): it contains a single captured packet, or a portion of it. It represents an evolution of the original Packet Block ([Section 4.5](#)).
- o Simple Packet Block ([Section 4.4](#)): it contains a single captured packet, or a portion of it, with only a minimal set of information about it.
- o Name Resolution Block ([Section 4.6](#)): it defines the mapping from numeric addresses present in the packet dump and the canonical name counterpart.
- o Interface Statistics Block ([Section 4.7](#)): it defines how to store some statistical data (e.g. packet dropped, etc) which can be useful to understand the conditions in which the capture has been made.

OBSOLETE blocks SHOULD NOT appear in newly written files (but left here for reference):

- o Packet Block ([Section 4.5](#)): it contains a single captured packet, or a portion of it. It should be considered OBSOLETE, and superseded by the Enhanced Packet Block ([Section 4.3](#)).

EXPERIMENTAL blocks are considered interesting but the authors believe that they deserve more in-depth discussion before being defined:

- o Alternative Packet Blocks
- o Compression Block
- o Encryption Block
- o Fixed Length Block
- o Directory Block
- o Traffic Statistics and Monitoring Blocks

- o Event/Security Blocks

3.3. Logical Block Hierarchy

The blocks build a logical hierarchy as they refer to each other. Figure 2 shows the logical hierarchy of the currently defined blocks in the form of a "tree view":

```

Section Header
|
+- Interface Description
|   +- Simple Packet
|   +- Enhanced Packet
|   +- Interface Statistics
|
+- Name Resolution

```

Figure 2: Logical block Hierarchy of a pcapng file.

For example: each captured packet refers to a specific capture interface, the interface itself refers to a specific section.

3.4. Physical File Layout

The file MUST begin with a Section Header Block. However, more than one Section Header Block can be present on the dump, each one covering the data following it till the next one (or the end of file). A Section includes the data delimited by two Section Header Blocks (or by a Section Header Block and the end of the file), including the first Section Header Block.

In case an application cannot read a Section because of different version number, it MUST skip everything until the next Section Header Block. Note that, in order to properly skip the blocks until the next section, all blocks MUST have the fields Type and Length at the beginning. This is a mandatory requirement that MUST be maintained in future versions of the block format.

Figure 3 shows a typical file configuration, with a single Section Header that covers the whole file.

```

+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| SHB v1.0 |                                     Data |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Figure 3: File structure example: Typical configuration with a single Section Header Block.

Figure 4 shows a file that contains three headers, and is normally the result of file concatenation. An application that understands only version 1.0 of the file format skips the intermediate section and restart processing the packets after the third Section Header.

```
|-- 1st Section  --|-- 2nd Section  --|-- 3rd Section  --|
|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| SHB v1.0 | Data | SHB V1.1 | Data | SHB V1.0 | Data |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

Figure 4: File structure example: three Section Header Blocks in a single file.

Figure 5 shows a file comparable to a "classic libpcap" file - the minimum for a useful capture file. It contains a single Section Header Block (SHB), a single Interface Description Block (IDB) and a few Enhanced Packet Blocks (EPB).

```
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| SHB | IDB | EPB | EPB | ... | EPB |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

Figure 5: File structure example: a pcapng file similar to a classical libpcap file.

Figure 6 shows a complex example file. In addition to the minimum file above, it contains packets captured from three interfaces, capturing on the third of which begins after packets have arrived on other interfaces, and also includes some Name Resolution Blocks (NRB) and an Interface Statistics Block (ISB).

```
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| SHB | IDB | IDB | EPB | NRB |...| IDB | EPB | ISB | NRB | EPB | EPB |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

Figure 6: File structure example: more complex pcapng file.

The last example should make it obvious that the block structure makes the file format very flexible compared to the classical libpcap format.

3.5. Options

All the block bodies have the possibility to embed optional fields. Optional fields can be used to insert some information that may be useful when reading data, but that is not really needed for packet

processing. Therefore, each tool can either read the content of the optional fields (if any), or skip some of them or even all at once.

Skipping all the optional fields at once is straightforward because most of the blocks are made of a first part with fixed format, and a second optional part. Therefore, the Block Length field (present in the General Block Structure, see [Section 3.1](#)) can be used to skip everything till the next block.

Options are a list of Type - Length - Value fields, each one containing a single value:

- o Option Type (2 bytes): it contains the code that specifies the type of the current TLV record. Option types whose Most Significant Bit is equal to one are reserved for local use; therefore, there is no guarantee that the code used is unique among all capture files (generated by other applications). In case of vendor-specific extensions that have to be identified uniquely, vendors MUST request an Option Code whose MSB is equal to zero.
- o Option Length (2 bytes): it contains the actual length of the following 'Option Value' field without the padding bytes.
- o Option Value (variable length): it contains the value of the given option, aligned to a 32-bit boundary. The actual length of this field (i.e. without the padding bytes) is specified by the Option Length field.

If an option's value is a string, the value is not necessarily zero-terminated. Software that reads these files MUST NOT assume that strings are zero-terminated, and MUST treat a zero byte as a string terminator.

Options may be repeated several times (e.g. an interface that has several IP addresses associated to it) TODO: mention for each option, if it can/shouldn't appear more than one time. The option list is terminated by a Option which uses the special 'End of Option' code (opt_endofopt).

The format of the optional fields is shown in Figure 7.

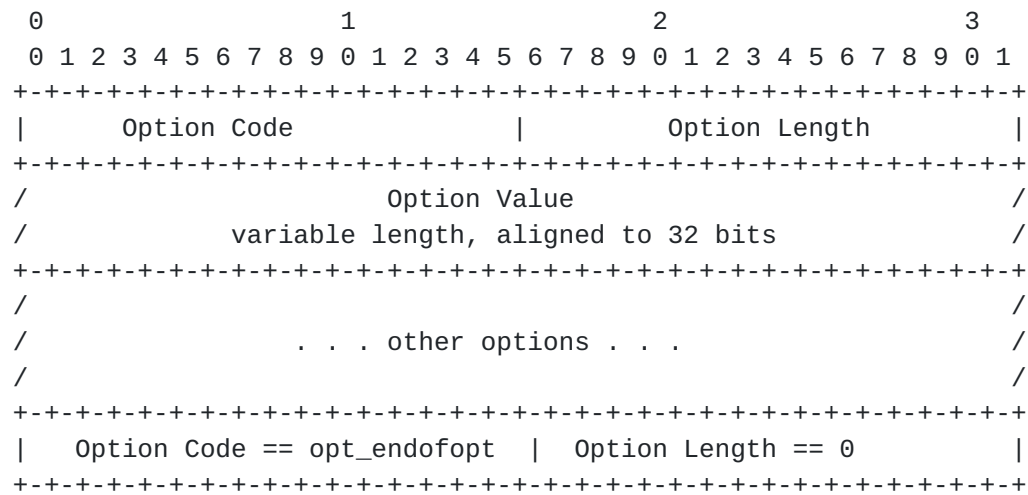


Figure 7: Options format.

The following codes can always be present in any optional field:

Name	Code	Length	Description	Example(s)
opt_endofopt	0	0	It delimits the end of the optional fields. This block cannot be repeated within a given list of options.	
opt_comment	1	variable	A UTF-8 string containing a comment that is associated to the current block.	"This packet is the beginning of all of our problems" / "Packets 17-23 showing a bogus TCP retransmission, as reported in bugzilla entry 1486!" / "Captured at the southern plant" / "I've checked again, now it's working ok" / ...

3.6. Data format

Endianess

Data contained in each section will always be saved according to the characteristics (little endian / big endian) of the dumping machine. This refers to all the fields that are saved as numbers and that span over two or more bytes.

The approach of having each section saved in the native format of the generating host is more efficient because it avoids translation of data when reading / writing on the host itself, which is the most common case when generating/processing capture dumps.

Please note: The endianness is indicated by the Section Header Block ([Section 4.1](#)). As this block can appear several times in a pcapng file, a single file can contain both endianness variants!

Alignment

All fields of this specification uses proper alignment for 16- and 32-bit values. This makes it easier and faster to read/write file contents if using techniques like memory mapped files.

The alignment bytes (marked in this document e.g. with "aligned to 32 bits") MUST be filled with zero bytes.

Please note: 64-bit values are not aligned to 64-bit boundaries. This is because the file is naturally aligned to 32-bit boundaries only. Special care MUST be taken when reading and writing such values. TODO: the spec is not too consistent wrt how 64-bit values are saved. in the Packet blocks we clearly specify where the low and high 32-bits of a 64-bit timestamp should be saved. In the SHB we do use the endianness of the machine when we save the section length.

[4.](#) Block Definition

This section details the format of the body of the blocks currently defined.

[4.1.](#) Section Header Block (mandatory)

The Section Header Block is mandatory. It identifies the beginning of a section of the capture dump file. The Section Header Block does not contain data but it rather identifies a list of blocks (interfaces, packets) that are logically correlated. Its format is shown in Figure 8.

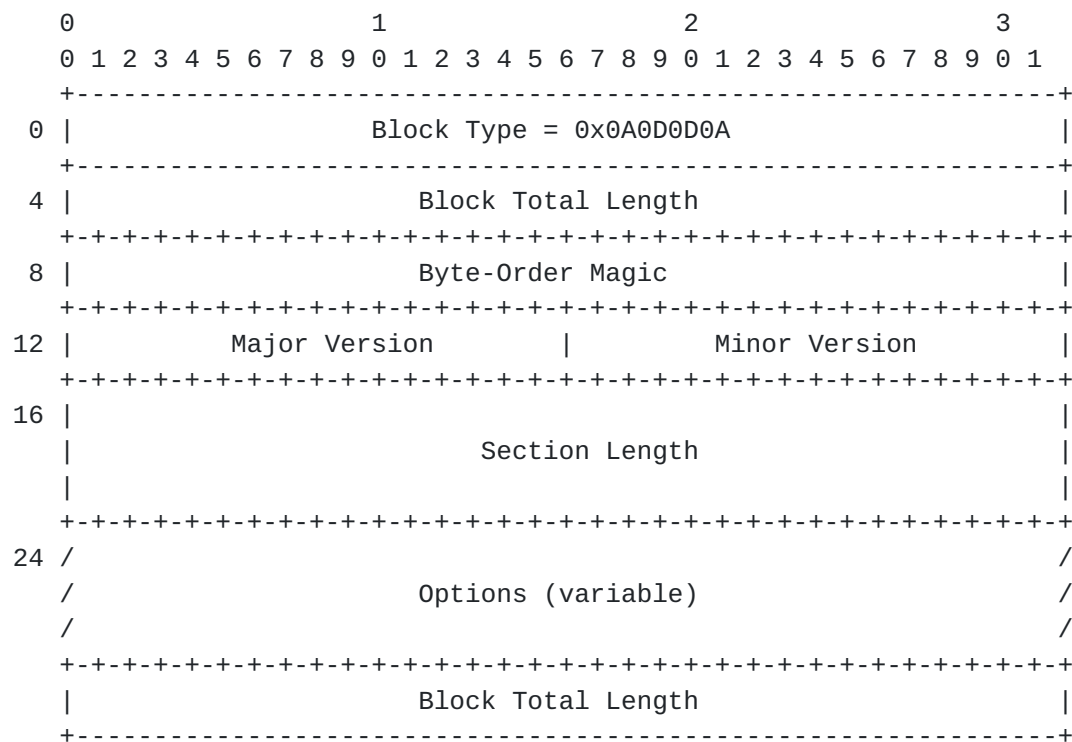


Figure 8: Section Header Block format.

The meaning of the fields is:

- o Block Type: The block type of the Section Header Block is the integer corresponding to the 4-char string "\r\n\n\r" (0x0A0D0D0A). This particular value is used for 2 reasons:
 1. This number is used to detect if a file has been transferred via FTP or HTTP from a machine to another with an inappropriate ASCII conversion. In this case, the value of this field will differ from the standard one ("\r\n\n\r") and the reader can detect a possibly corrupted file.
 2. This value is palindromic, so that the reader is able to recognize the Section Header Block regardless of the endianness of the section. The endianness is recognized by reading the Byte Order Magic, that is located 8 bytes after the Block Type.
- o Block Total Length: total size of this block, as described in [Section 3.1](#).
- o Byte-Order Magic: magic number, whose value is the hexadecimal number 0x1A2B3C4D. This number can be used to distinguish

sections that have been saved on little-endian machines from the ones saved on big-endian machines.

- o Major Version: number of the current mayor version of the format. Current value is 1. This value should change if the format changes in such a way that tools that can read the new format could not read the old format (i.e., the code would have to check the version number to be able to read both formats).
- o Minor Version: number of the current minor version of the format. Current value is 0. This value should change if the format changes in such a way that tools that can read the new format can still automatically read the new format but code that can only read the old format cannot read the new format.
- o Section Length: 64-bit value specifying the length in bytes of the following section, excluding the Section Header Block itself. This field can be used to skip the section, for faster navigation inside large files. Section Length equal -1 (0xFFFFFFFFFFFFFFF) means that the size of the section is not specified, and the only way to skip the section is to parse the blocks that it contains. Please note that if this field is valid (i.e. not -1), its value is always aligned to 32 bits, as all the blocks are aligned to 32-bit boundaries. Also, special care should be taken in accessing this field: since the alignment of all the blocks in the file is 32-bit, this field is not guaranteed to be aligned to a 64-bit boundary. This could be a problem on 64-bit workstations.
- o Options: optionally, a list of options (formatted according to the rules defined in [Section 3.5](#)) can be present.

Adding new block types or options would not necessarily require that either Major or Minor numbers be changed, as code that does not know about the block type or option could just skip it; only if skipping a block or option does not work should the minor version number be changed.

Aside from the options defined in [Section 3.5](#), the following options are valid within this block:

Name	Code	Length	Description	Example(s)
shb_hardware	2	variable	An UTF-8 string containing the description of the hardware used to create this section.	"x86 Personal Computer" / "Sun Sparc Workstation" / ...
shb_os	3	variable	An UTF-8 string containing the name of the operating system used to create this section.	"Windows XP SP2" / "openSUSE 10.2" / ...
shb_userappl	4	variable	An UTF-8 string containing the name of the application used to create this section.	"dumpcap V0.99.7" / ...

4.2. Interface Description Block (mandatory)

The Interface Description Block is mandatory. This block is needed to specify the characteristics of the network interface on which the capture has been made. In order to properly associate the captured data to the corresponding interface, the Interface Description Block MUST be defined before any other block that uses it; therefore, this block is usually placed immediately after the Section Header Block. However, if capturing on an interface starts in the middle of a capture, an Interface Description Block can appear after other blocks, including blocks for packets on interfaces for which Interface Description Blocks have already appeared.

An Interface Description Block is valid only inside the section which it belongs to. The structure of a Interface Description Block is shown in Figure 9.

In addition to the options defined in [Section 3.5](#), the following options are valid within this block:

Name	Code	Length	Description	Example(s)
if_name	2	Variable	A UTF-8 string containing the name of the device used to capture data.	"eth0" / "\\Device\\NPF_{AD1CE675-96D0-47C5-ADD0-2504B9126B68}" / ...
if_description	3	Variable	A UTF-8 string containing the description of the device used to capture data.	"Broadcom NetXtreme" / "First Ethernet Interface" / ...
if_ipv4address	4	8	Interface network address and netmask. This option can be repeated multiple times within the same Interface Description Block when multiple IPv4 addresses are assigned	192 168 1 1 255 255 255 0

			to the in	
			terface.	
if_IPv6a	5	17	Interface	2001:0db8:85a3:08d3:1319:8a2e
ddr			network	:0370:7344/64 is written (in
			address	hex) as "20 01 0d b8 85 a3 08
			and	d3 13 19 8a 2e 03 70 73 44
			prefix	40"
			length	
			(stored	
			in the	
			last	
			byte).	
			This	
			option	
			can be	
			repeated	
			multiple	
			times	
			within	
			the same	
			Interface	
			Descripti	
			on Block	
			when	
			multiple	
			IPv6	
			addresses	
			are	
			assigned	
			to the in	
			terface.	
if_MACad	6	6	Interface	00 01 02 03 04 05
dr			Hardware	
			MAC	
			address	
			(48	
			bits).	
if_EUIad	7	8	Interface	02 34 56 FF FE 78 9A BC
dr			Hardware	
			EUI	
			address	
			(64	
			bits), if	
			available	
			.	
if_speed	8	8	Interface	100000000 for 100Mbps
			speed (in	
			bps).	

if_tsres	9	1	Resolution of time stamps.	6
ol			If the Most Significant Bit is equal to zero, the remaining bits indicates the resolution of the timestamp as a negative power of 10 (e.g. 6 means milliseconds resolution, time stamps are the number of microseconds since 1/1/1970). If the Most Significant Bit is equal to one, the remaining bits indicates the resolution as a negative power of 2 (e.g. 10 means 1/1024 of second).	

			If this option is not present, a resolution of 10^{-6} is assumed (i.e. timestamps have the same resolution of the standard 'libpcap' timestamps).	
if_tzone	10	4	Time zone for GMT support (TODO: specify better).	TODO: give a good example
if_filter	11	variable	The filter (e.g. "capture only TCP traffic") used to capture traffic. The first byte of the Option Data keeps a code of the filter used (e.g. if this is a libpcap string, or BPF	00 "tcp port 23 and host 192.0.2.5"

			bytecode,	
			and	
			more).	
			More	
			details	
			about	
			this	
			format	
			will be	
			presented	
			in	
			Appendix	
			XXX	
			(TODO).	
			(TODO:	
			better	
			use	
			different	
			options	
			for	
			different	
			fields?	
			e.g. if_f	
			ilter_pca	
			p, if_fil	
			ter_bpf,	
			...)	
if_os	12	variable	A UTF-8	"Windows XP SP2" / "openSUSE
			string co	10.2" / ...
			ntaining	
			the name	
			of the	
			operating	
			system of	
			the	
			machine	
			in which	
			this	
			interface	
			is instal	
			led. This	
			can be	
			different	
			from the	
			same info	
			rmation	
			that can	
			be	

			contained	
			by the	
			Section	
			Header	
			Block	
			(Section	
			4.1)	
			because	
			the	
			capture	
			can have	
			been done	
			on a	
			remote	
			machine.	
if_fcsle	13	1	An	4
n			integer	
			value	
			that	
			specified	
			the	
			length of	
			the Frame	
			Check	
			Sequence	
			(in bits)	
			for this	
			interface	
			. For	
			link	
			layers	
			whose FCS	
			length	
			can	
			change	
			during	
			time, the	
			Packet	
			Block	
			Flags	
			Word can	
			be used	
			(see	
			Appendix	
			A).	
if_tsoff	14	8	A 64-bit	1234
set			integer	
			value	

that specifies an offset (in seconds) that must be added to the timestamp of each packet to obtain the absolute timestamp of a packet. If the option is missing, the timestamps stored in the packet MUST be considered absolute timestamps. The time zone of the offset can be specified with the option `if_tzone`.
TODO: won't a `if_tsoffset_low` for fractional second offsets be useful for highly sy

				ncronized	
				capture	
				systems?	
+-----+	+-----+	+-----+	+-----+	+-----+	+-----+

4.3. Enhanced Packet Block (optional)

An Enhanced Packet Block is the standard container for storing the packets coming from the network. The Enhanced Packet Block is optional because packets can be stored either by means of this block or the Simple Packet Block, which can be used to speed up dump generation. The format of an Enhanced Packet Block is shown in Figure 10.

The Enhanced Packet Block is an improvement over the original Packet Block ([Section 4.5](#)):

- o it stores the Interface Identifier as a 32-bit integer value. This is a requirement when a capture stores packets coming from a large number of interfaces
- o differently from the Packet Block ([Section 4.5](#)), the number of packets dropped by the capture system between this packet and the previous one is not stored in the header, but rather in an option of the block itself.

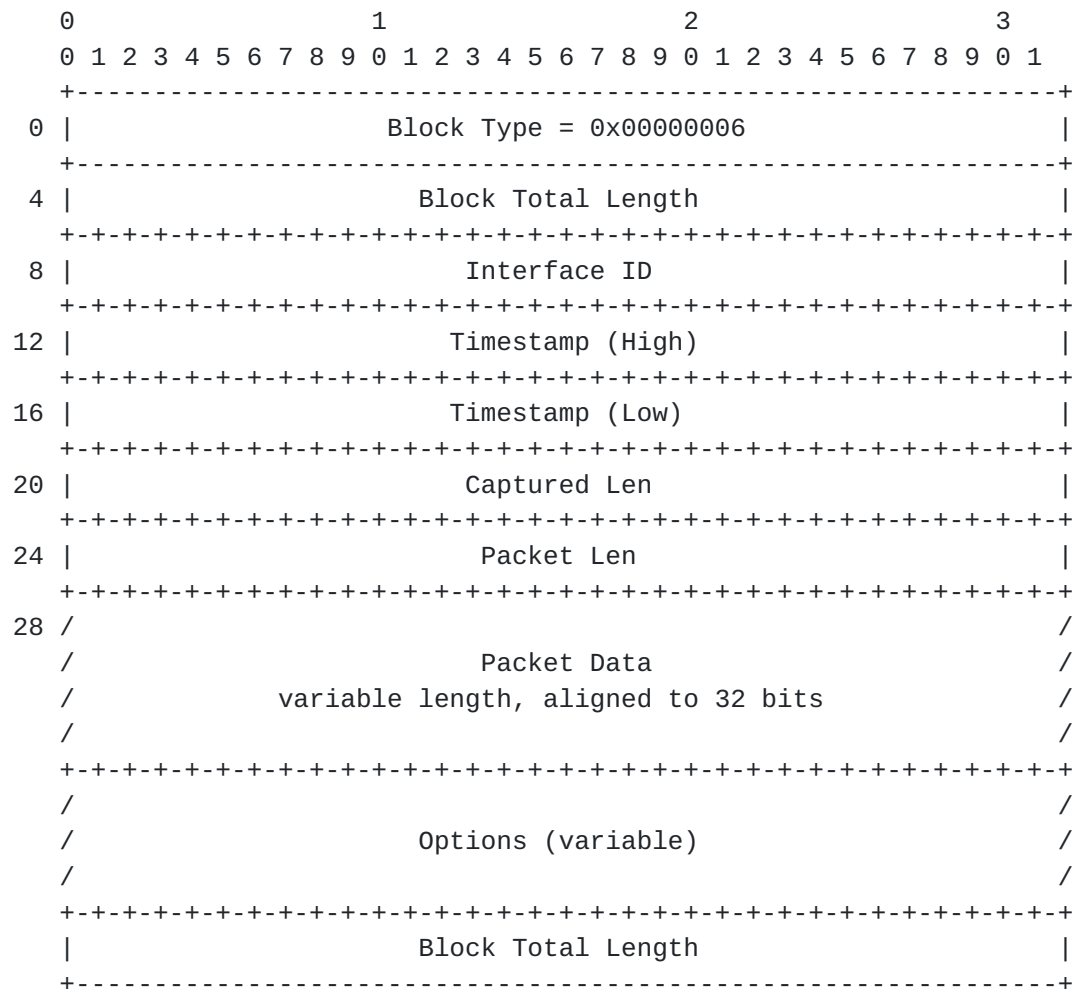


Figure 10: Enhanced Packet Block format.

The Enhanced Packet Block has the following fields:

- o Block Type: The block type of the Enhanced Packet Block is 6.
- o Block Total Length: total size of this block, as described in [Section 3.1](#).
- o Interface ID: it specifies the interface this packet comes from; the correct interface will be the one whose Interface Description Block (within the current Section of the file) is identified by the same number (see [Section 4.2](#)) of this field. The interface ID MUST be valid, which means that an matching interface description block MUST exist.
- o Timestamp (High) and Timestamp (Low): high and low 32-bits of a 64-bit quantity representing the timestamp. The timestamp is a single 64-bit unsigned integer representing the number of units

since 1/1/1970 00:00:00 UTC. The way to interpret this field is specified by the 'if_tsresol' option (see Figure 9) of the Interface Description block referenced by this packet. Please note that differently from the libpcap file format, timestamps are not saved as two 32-bit values accounting for the seconds and microseconds since 1/1/1970. They are saved as a single 64-bit quantity saved as two 32-bit words.

- o Captured Len: number of bytes captured from the packet (i.e. the length of the Packet Data field). It will be the minimum value among the actual Packet Length and the snapshot length (defined in Figure 9). The value of this field does not include the padding bytes added at the end of the Packet Data field to align the Packet Data Field to a 32-bit boundary.
- o Packet Len: actual length of the packet when it was transmitted on the network. It can be different from Captured Len if the user wants only a snapshot of the packet.
- o Packet Data: the data coming from the network, including link-layer headers. The actual length of this field is Captured Len. The format of the link-layer headers depends on the LinkType field specified in the Interface Description Block (see [Section 4.2](#)) and it is specified in the entry for that format in the the tcpdump.org link-layer header types registry [2].
- o Options: optionally, a list of options (formatted according to the rules defined in [Section 3.5](#)) can be present.

In addition to the options defined in [Section 3.5](#) and in the Packet Block (Table 1), the following options are valid within this block:

Name	Code	Length	Description	Example(s)
epb_flags	2	4	A flags word containing link-layer information. A complete specification of the allowed flags can be found in Appendix A .	0
epb_hash	3	variable	This option contains a hash of the packet. The first byte specifies the	2 EC 1D 87 97 / 3 45 6e c2 17 7c 10 1e 3c 2e 99

			hashing algorithm,	6e c2 9a	
			while the	3d 50 8e	
			following bytes		
			contain the actual		
			hash, whose size		
			depends on the		
			hashing algorithm,		
			and hence from the		
			value in the first		
			byte. The hashing		
			algorithm can be:		
			2s complement		
			(algorithm byte =		
			0, size=XXX), XOR		
			(algorithm byte =		
			1, size=XXX),		
			CRC32 (algorithm		
			byte = 2, size =		
			4), MD-5		
			(algorithm byte =		
			3, size=XXX),		
			SHA-1 (algorithm		
			byte = 4,		
			size=XXX). The		
			hash covers only		
			the packet, not		
			the header added		
			by the capture		
			driver: this gives		
			the possibility to		
			calculate it		
			inside the network		
			card. The hash		
			allows easier		
			comparison/merging		
			of different		
			capture files, and		
			reliable data		
			transfer between		
			the data		
			acquisition system		
			and the capture		
			library. (TODO:		
			the text above		
			uses "first bit",		
			but shouldn't this		
			be "first		
			byte"?!?)		

epb_dropcount	4	8	A 64-bit integer	0	
			value specifying		
			the number of		
			packets lost (by		
			the interface and		
			the operating		
			system) between		
			this packet and		
			the preceding one.		
+-----+	+-----+	+-----+	+-----+	+-----+	+-----+

4.4. Simple Packet Block (optional)

The Simple Packet Block is a lightweight container for storing the packets coming from the network. Its presence is optional.

A Simple Packet Block is similar to a Packet Block (see [Section 4.5](#)), but it is smaller, simpler to process and contains only a minimal set of information. This block is preferred to the standard Packet Block when performance or space occupation are critical factors, such as in sustained traffic dump applications. A capture file can contain both Packet Blocks and Simple Packet Blocks: for example, a capture tool could switch from Packet Blocks to Simple Packet Blocks when the hardware resources become critical.

The Simple Packet Block does not contain the Interface ID field. Therefore, it MUST be assumed that all the Simple Packet Blocks have been captured on the interface previously specified in the first Interface Description Block.

Figure 11 shows the format of the Simple Packet Block.

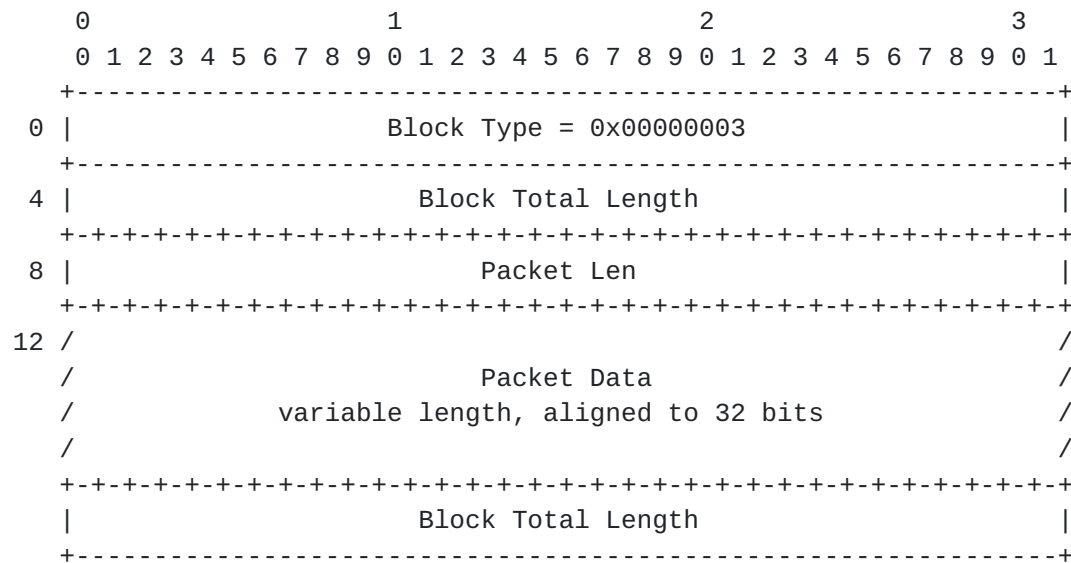


Figure 11: Simple Packet Block format.

The Simple Packet Block has the following fields:

- o Block Type: The block type of the Simple Packet Block is 3.
- o Block Total Length: total size of this block, as described in [Section 3.1](#).
- o Packet Len: actual length of the packet when it was transmitted on the network. Can be different from captured len if the packet has been truncated by the capture process.
- o Packet Data: the data coming from the network, including link-layer headers. The length of this field can be derived from the field Block Total Length, present in the Block Header, and it is the minimum value among the SnapLen (present in the Interface Description Block) and the Packet Len (present in this header). The format of the link-layer headers depends on the LinkType field specified in the Interface Description Block (see [Section 4.2](#)) and it is specified in the entry for that format in the the tcpdump.org link-layer header types registry [3].

The Simple Packet Block does not contain the timestamp because this is often one of the most costly operations on PCs. Additionally, there are applications that do not require it; e.g. an Intrusion Detection System is interested in packets, not in their timestamp.

A Simple Packet Block cannot be present in a Section that has more than one interface because of the impossibility to refer to the correct one (it does not contain any Interface ID field).

The Simple Packet Block is very efficient in term of disk space: a snapshot whose length is 100 bytes requires only 16 bytes of overhead, which corresponds to an efficiency of more than 86%.

4.5. Packet Block (obsolete!)

The Packet Block is marked obsolete, better use the Enhanced Packet Block instead!

A Packet Block is the standard container for storing the packets coming from the network. The Packet Block is optional because packets can be stored either by means of this block or the Simple Packet Block, which can be used to speed up dump generation. The format of a packet block is shown in Figure 12.

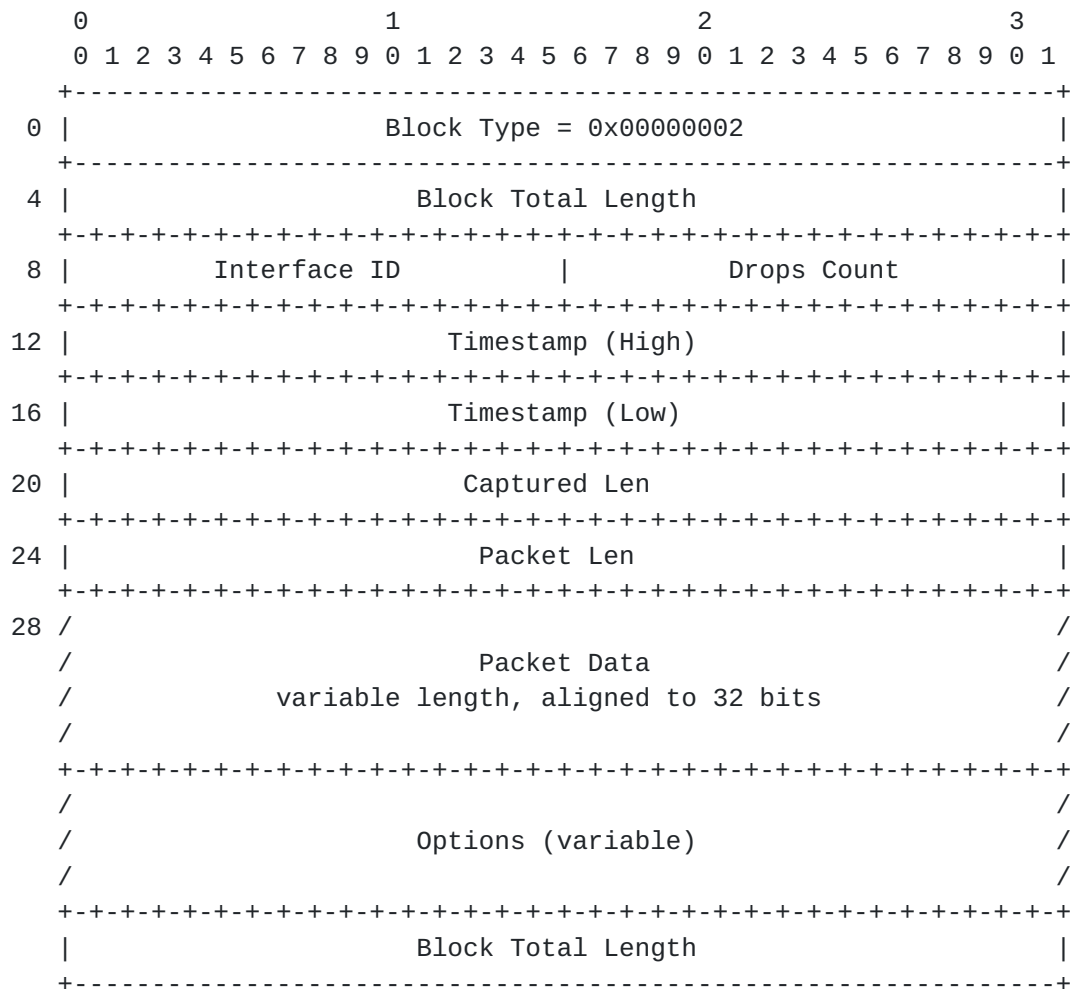


Figure 12: Packet Block format.

The Packet Block has the following fields:

- o Block Type: The block type of the Packet Block is 2.
- o Block Total Length: total size of this block, as described in [Section 3.1](#).
- o Interface ID: it specifies the interface this packet comes from; the correct interface will be the one whose Interface Description Block (within the current Section of the file) is identified by the same number (see [Section 4.2](#)) of this field. The interface ID MUST be valid, which means that an matching interface description block MUST exist.
- o Drops Count: a local drop counter. It specifies the number of packets lost (by the interface and the operating system) between this packet and the preceding one. The value xFFFF (in hexadecimal) is reserved for those systems in which this information is not available.
- o Timestamp (High) and Timestamp (Low): timestamp of the packet. The format of the timestamp is the same already defined in the Enhanced Packet Block ([Section 4.3](#)).
- o Captured Len: number of bytes captured from the packet (i.e. the length of the Packet Data field). It will be the minimum value among the actual Packet Length and the snapshot length (SnapLen defined in Figure 9). The value of this field does not include the padding bytes added at the end of the Packet Data field to align the Packet Data Field to a 32-bit boundary.
- o Packet Len: actual length of the packet when it was transmitted on the network. Can be different from Captured Len if the user wants only a snapshot of the packet.
- o Packet Data: the data coming from the network, including link-layer headers. The actual length of this field is Captured Len. The format of the link-layer headers depends on the LinkType field specified in the Interface Description Block (see [Section 4.2](#)) and it is specified in the entry for that format in the the tcpdump.org link-layer header types registry [4].
- o Options: optionally, a list of options (formatted according to the rules defined in [Section 3.5](#)) can be present.

In addition to the options defined in [Section 3.5](#), the following options are valid within this block:

Name	Code	Length	Description	Example(s)
pack_flags	2	4	Same as epb_flags of the enhanced packet block.	0
pack_hash	3	variable	Same as epb_hash of the enhanced packet block.	2 EC 1D 87 97 / 3 45 6e c2 17 7c 10 1e 3c 2e 99 6e c2 9a 3d 50 8e

Table 1

4.6. Name Resolution Block (optional)

The Name Resolution Block is used to support the correlation of numeric addresses (present in the captured packets) and their corresponding canonical names and it is optional. Having the literal names saved in the file, this prevents the need of a name resolution in a delayed time, when the association between names and addresses can be different from the one in use at capture time. Moreover, the Name Resolution Block avoids the need of issuing a lot of DNS requests every time the trace capture is opened, and allows to have name resolution also when reading the capture with a machine not connected to the network.

A Name Resolution Block is normally placed at the beginning of the file, but no assumptions can be taken about its position. Name Resolution Blocks can be added in a second time by tools that process the file, like network analyzers.

The format of the Name Resolution Block is shown in Figure 13.

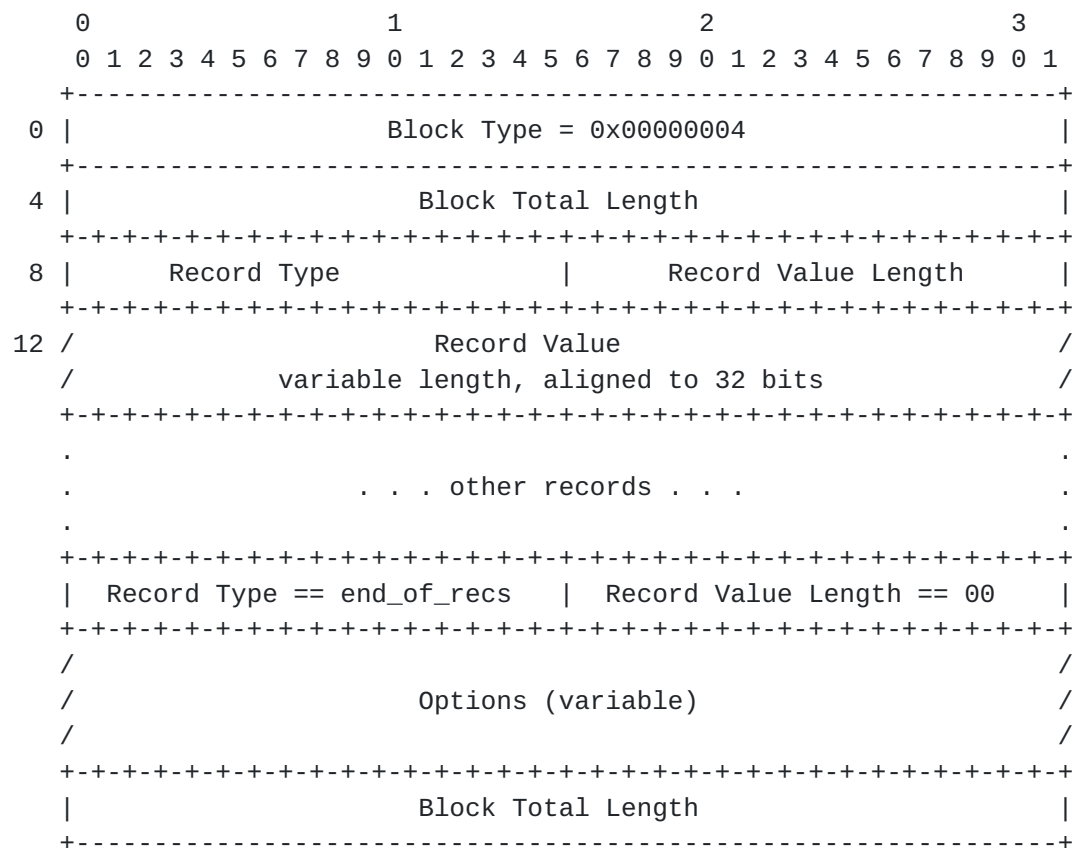


Figure 13: Name Resolution Block format.

The Name Resolution Block has the following fields:

- o Block Type: The block type of the Name Resolution Block is 4.
- o Block Total Length: total size of this block, as described in [Section 3.1](#).

This is followed by a zero-terminated list of records (in the TLV format), each of which contains an association between a network address and a name. There are three possible types of records:

Name	Code	Length	Description	Example(s)
nres_endofrecord	0	0	It delimits the end of name resolution records. This record is needed to determine when the list of name resolution records has ended and some options (if any) begin.	
nres_ip4record	1	Variable	Specifies an IPv4 address (contained in the first 4 bytes), followed by one or more zero-terminated strings containing the DNS entries for that address.	127 0 0 1 "localhost"
nres_ip6record	2	Variable	Specifies an IPv6 address (contained in the first 16 bytes), followed by one or more zero-terminated strings containing the DNS entries for that address.	20 01 0d b8 00 00 00 00 00 00 00 00 12 34 56 78 "somehost"

Table 2

Each Record Value is aligned to a 32-bit boundary. The corresponding Record Value Length reflects the actual length of the Record Value; it does not include the lengths of the Record Type or the Record Value Length.

After the list of Name Resolution Records, optionally, a list of options (formatted according to the rules defined in [Section 3.5](#)) can be present.

In addition to the options defined in [Section 3.5](#), the following options are valid within this block:

Name	Code	Length	Description	Example(s)
ns_dnsname	2	Variable	A UTF-8 string containing the name of the machine (DNS server) used to perform the name resolution.	"our_nameserver"
ns_dnsIP4addr	3	4	The IPv4 address of the DNS server.	192 168 0 1
ns_dnsIP6addr	4	16	The IPv6 address of the DNS server.	20 01 0d b8 00 00 00 00 00 00 12 34 56 78

[4.7.](#) Interface Statistics Block (optional)

The Interface Statistics Block contains the capture statistics for a given interface and it is optional. The statistics are referred to the interface defined in the current Section identified by the Interface ID field. An Interface Statistics Block is normally placed at the end of the file, but no assumptions can be taken about its position - it can even appear multiple times for the same interface.

The format of the Interface Statistics Block is shown in Figure 14.

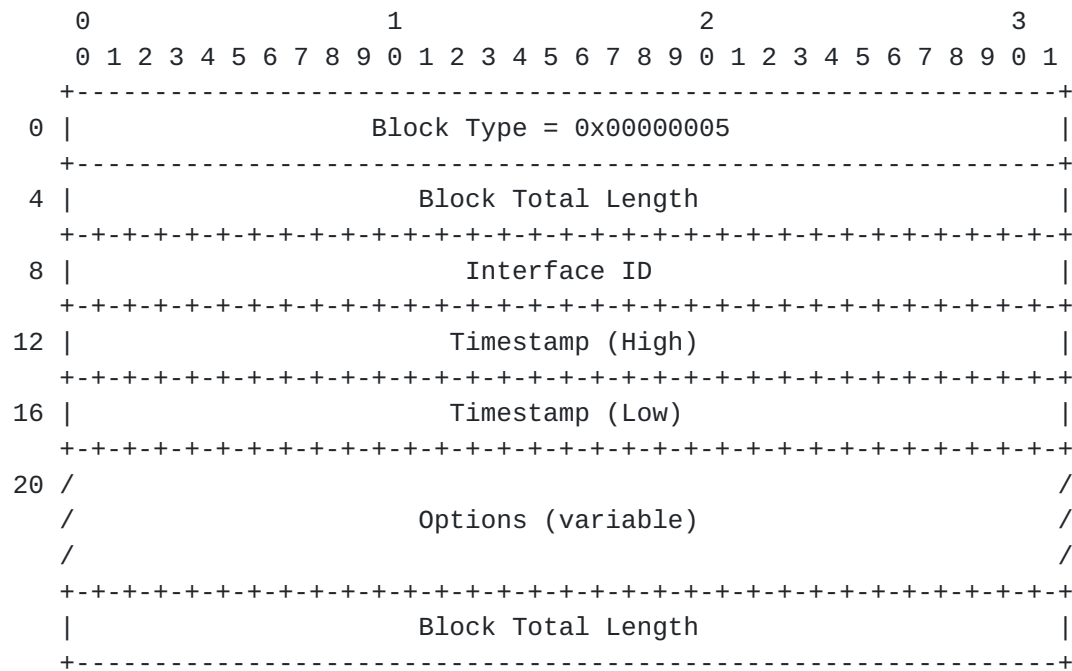


Figure 14: Interface Statistics Block format.

The fields have the following meaning:

- o Block Type: The block type of the Interface Statistics Block is 5.
- o Block Total Length: total size of this block, as described in [Section 3.1](#).
- o Interface ID: it specifies the interface these statistics refers to; the correct interface will be the one whose Interface Description Block (within the current Section of the file) is identified by same number (see [Section 4.2](#)) of this field. Please note: in former versions of this document, this field was 16 bits only. As this differs from its usage in other places of this doc and as this block was not used "in the wild" before (as to the knowledge of the authors), it seems reasonable to change it to 32 bits!
- o Timestamp: time this statistics refers to. The format of the timestamp is the same already defined in the Enhanced Packet Block ([Section 4.3](#)).
- o Options: optionally, a list of options (formatted according to the rules defined in [Section 3.5](#)) can be present.

All the statistic fields are defined as options in order to deal with systems that do not have a complete set of statistics. Therefore, In

addition to the options defined in [Section 3.5](#), the following options are valid within this block:

Name	Code	Length	Description	
isb_starttime	2	8	Time in which the capture started; time will be stored in two blocks of four bytes each. The format of the timestamp is the same already defined in the Enhanced Packet Block (Section 4.3).	97 c3 04 00 aa 47 ca 64 (Little Endian, decodes to 06/29/2012 06:16:50 UTC)
isb_endtime	3	8	Time in which the capture ended; ; time will be stored in two blocks of four bytes each. The format of the timestamp is the same already defined in the Enhanced Packet Block (Section 4.3).	96 c3 04 00 73 89 6a 65 (Little Endian, decodes to 06/29/2012 06:17:00 UTC)
isb_ifrecv	4	8	Number of packets received from the physical interface starting from the beginning of the capture.	100
isb_ifdrop	5	8	Number of packets dropped by the interface due to lack of resources starting from the beginning of the capture.	0
isb_filteraccept	6	8	Number of packets accepted by	100

			filter starting	
			from the	
			beginning of the	
			capture.	
isb_osdrop	7	8	Number of packets	0
			dropped by the	
			operating system	
			starting from the	
			beginning of the	
			capture.	
isb_usrdeliv	8	8	Number of packets	0
			delivered to the	
			user starting	
			from the	
			beginning of the	
			capture. The	
			value contained	
			in this field can	
			be different from	
			the value	
			'isb_filteraccept	
			- isb_osdrop'	
			because some	
			packets could	
			still lay in the	
			OS buffers when	
			the capture	
			ended.	
+-----+-----+-----+-----+-----+				

All the fields that refer to packet counters are 64-bit values, represented with the byte order of the current section. Special care must be taken in accessing these fields: since all the blocks are aligned to a 32-bit boundary, such fields are not guaranteed to be aligned on a 64-bit boundary.

5. Experimental Blocks (deserved to a further investigation)

5.1. Alternative Packet Blocks (experimental)

Can some other packet blocks (besides the ones described in the previous paragraphs) be useful?

5.2. Compression Block (experimental)

The Compression Block is optional. A file can contain an arbitrary number of these blocks. A Compression Block, as the name says, is used to store compressed data. Its format is shown in Figure 15.

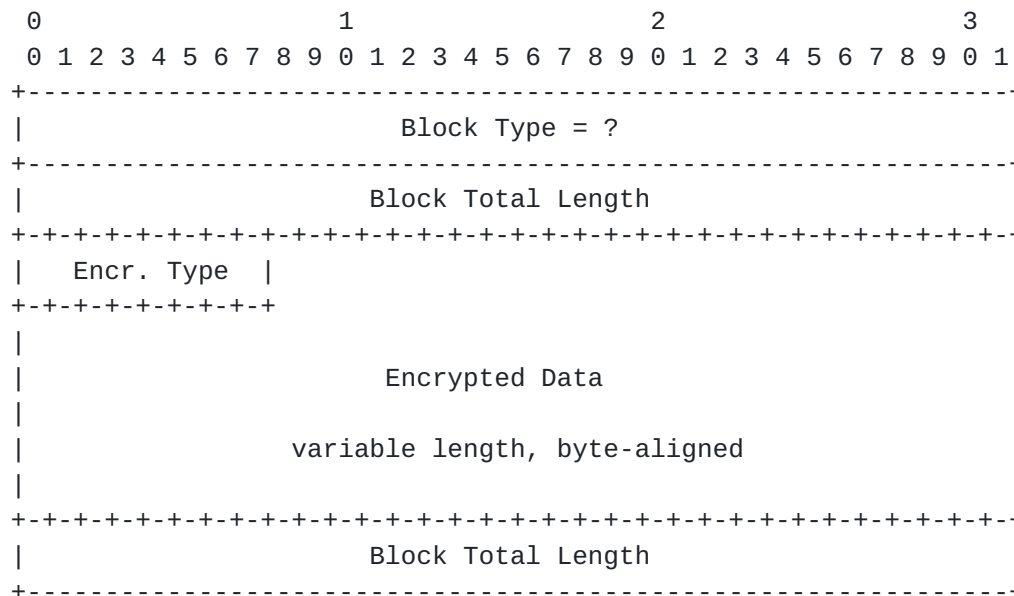


Figure 16: Encryption Block format.

The fields have the following meaning:

- o Block Type: The block type of the Encryption Block is not yet assigned.
- o Block Total Length: total size of this block, as described in [Section 3.1](#).
- o Encryption Type: specifies the encryption algorithm. Possible values for this field are ??? (TODO) NOTE: this block should probably contain other fields, depending on the encryption algorithm. To be defined precisely.
- o Encrypted Data: data of this block. Once decrypted, it originates other blocks.

5.4. Fixed Length Block (experimental)

The Fixed Length Block is optional. A file can contain an arbitrary number of these blocks. A Fixed Length Block can be used to optimize the access to the file. Its format is shown in Figure 17. A Fixed Length Block stores records with constant size. It contains a set of Blocks (normally Packet Blocks or Simple Packet Blocks), of which it specifies the size. Knowing this size a priori helps to scan the file and to load some portions of it without truncating a block, and is particularly useful with cell-based networks like ATM.

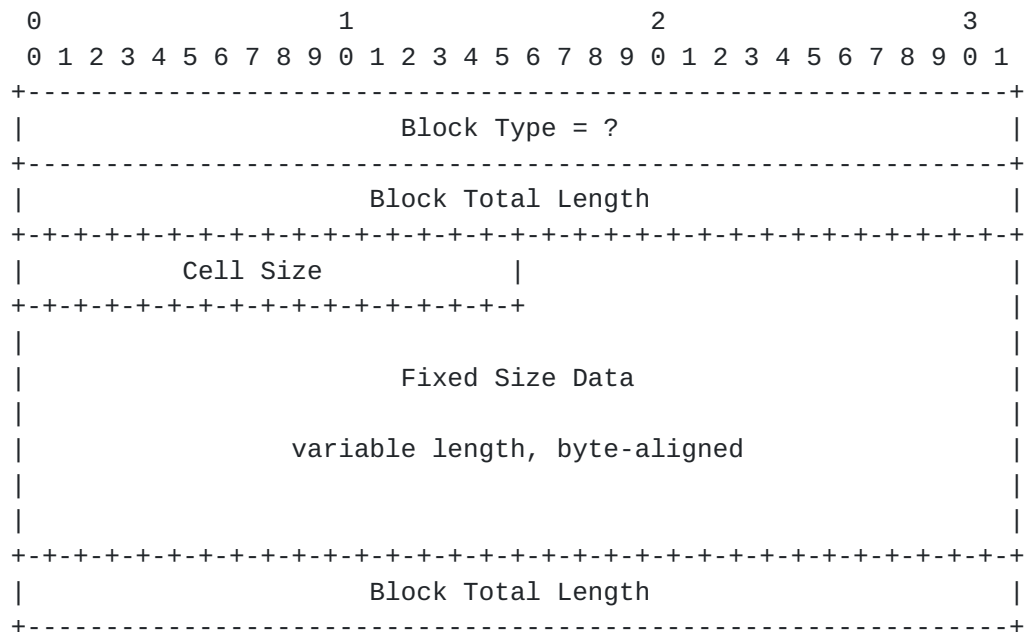


Figure 17: Fixed Length Block format.

The fields have the following meaning:

- o Block Type: The block type of the Fixed Length Block is not yet assigned.
- o Block Total Length: total size of this block, as described in [Section 3.1](#).
- o Cell size: the size of the blocks contained in the data field.
- o Fixed Size Data: data of this block.

5.5. Directory Block (experimental)

If present, this block contains the following information:

- o number of indexed packets (N)
- o table with position and length of any indexed packet (N entries)

A directory block **MUST** be followed by at least N packets, otherwise it **MUST** be considered invalid. It can be used to efficiently load portions of the file to memory and to support operations on memory mapped files. This block can be added by tools like network analyzers as a consequence of file processing.

5.6. Traffic Statistics and Monitoring Blocks (experimental)

One or more blocks could be defined to contain network statistics or traffic monitoring information. They could be used to store data collected from RMON or Netflow probes, or from other network monitoring tools.

5.7. Event/Security Block (experimental)

This block could be used to store events. Events could contain generic information (for example network load over 50%, server down...) or security alerts. An event could be:

- o skipped, if the application doesn't know how to do with it
- o processed independently by the packets. In other words, the applications skip the packets and process only the alerts
- o processed in relation to packets: for example, a security tool could load only the packets of the file that are near a security alert; a monitoring tool could skip the packets captured while the server was down.

6. Recommended File Name Extension: .pcapng

The recommended file name extension for the "PCAP Next Generation Dump File Format" specified in this document is ".pcapng".

On Windows and OS X, files are distinguished by an extension to their filename. Such an extension is technically not actually required, as applications should be able to automatically detect the pcapng file format through the "magic bytes" at the beginning of the file, as some other UNIX desktop environments do. However, using name extensions makes it easier to work with files (e.g. visually distinguish file formats) so it is recommended - though not required - to use .pcapng as the name extension for files following this specification.

Please note: To avoid confusion (like the current usage of .cap for a plethora of different capture file formats) other file name extensions than .pcapng should be avoided!

7. How to add Vendor / Domain specific extensions

TODO - explain the preferred way to add new block types and new options for existing blocks in more detail.

8. Conclusions

The file format proposed in this document should be very versatile and satisfy a wide range of applications. In the simplest case, it can contain a raw dump of the network data, made of a series of Simple Packet Blocks. In the most complex case, it can be used as a repository for heterogeneous information. In every case, the file remains easy to parse and an application can always skip the data it is not interested in; at the same time, different applications can share the file, and each of them can benefit of the information produced by the others. Two or more files can be concatenated obtaining another valid file.

9. Security Considerations

TBD.

10. IANA Considerations

TBD.

11. Acknowledgments

Loris Degioanni and Gianluca Varenni were coauthoring this document before it was submitted to the IETF.

The authors wish to thank Anders Broman, Ulf Lamping, Richard Sharpe and many others for their invaluable comments.

12. References

12.1. Normative References

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

12.2. URIs

[1] <http://www.tcpdump.org/linktypes.html>

[2] <http://www.tcpdump.org/linktypes.html>

[3] <http://www.tcpdump.org/linktypes.html>

[4] <http://www.tcpdump.org/linktypes.html>

Appendix A. Packet Block Flags Word

The Packet Block Flags Word is a 32-bit value that contains link-layer information about the packet.

The meaning of the bits is the following:

Bit Number	Description
0-1	Inbound / Outbound packet (00 = information not available, 01 = inbound, 10 = outbound)
2-4	Reception type (000 = not specified, 001 = unicast, 010 = multicast, 011 = broadcast, 100 = promiscuous).
5-8	FCS length, in bytes (0000 if this information is not available). This value overrides the if_fcslen option of the Interface Description Block, and is used with those link layers (e.g. PPP) where the length of the FCS can change during time.
9-15	Reserved (MUST be set to zero).
16-31	link-layer-dependent errors (Bit 31 = symbol error, Bit 30 = preamble error, Bit 29 = Start Frame Delimiter error, Bit 28 = unaligned frame error, Bit 27 = wrong Inter Frame Gap error, Bit 26 = packet too short error, Bit 25 = packet too long error, Bit 24 = CRC error, other?? are 16 bit enough?).

Appendix B. Standardized Block Type Codes

Every Block is uniquely identified by a 32-bit integer value, stored in the Block Header.

As pointed out in [Section 3.1](#), Block Types codes whose Most Significant Bit (bit 31) is set to 1 are reserved for local use by the application.

All the remaining Block Type codes (0x00000000 to 0x7FFFFFFF) are standardized by this document. A request should be sent to the authors of this document to add a new Standard Block Type code to the specification.

Here is a list of the Standardized Block Type Codes.

Block Type Code	Description
0x00000000	Reserved ???
0x00000001	Interface Description Block (Section 4.2)
0x00000002	Packet Block (Section 4.5)
0x00000003	Simple Packet Block (Section 4.4)
0x00000004	Name Resolution Block (Section 4.6)
0x00000005	Interface Statistics Block (Section 4.7)
0x00000006	Enhanced Packet Block (Section 4.3)
0x00000007	IRIG Timestamp Block (requested by Gianluca Varenni <gianluca.varenni@cacetech.com>, CACE Technologies LLC)
0x00000008	Arinc 429 in AFDX Encapsulation Information Block (requested by Gianluca Varenni <gianluca.varenni@cacetech.com>, CACE Technologies LLC)
0x0A0D0D0A	Section Header Block (Section 4.1)
0x0A0D0A00-0x0A0D0AFF	Reserved. Used to detect trace files corrupted because of file transfers using the HTTP protocol in text mode.
0x000A0D0A-0xFF0A0D0A	Reserved. Used to detect trace files corrupted because of file transfers using the HTTP protocol in text mode.
0x000A0D0D-0xFF0A0D0D	Reserved. Used to detect trace files corrupted because of file transfers using the HTTP protocol in text mode.
0x0D0D0A00-0x0D0D0AFF	Reserved. Used to detect trace files corrupted because of file transfers using the FTP protocol in text mode.

Authors' Addresses

Michael Tuexen (editor)
Muenster University of Applied Sciences
Stegerwaldstrasse 39
Steinfurt 48565
DE

Email: tuexen@fh-muenster.de

Fulvio Riso
Politecnico di Torino
Corso Duca degli Abruzzi, 24
Torino 10129
IT

Email: fulvio.riso@polito.it

Jasper Bongertz
Airbus Defence and Space CyberSecurity
Kanzlei 63c
Meerbusch 40667
DE

Email: jasper@packet-foo.com

Guy Harris

Email: guy@alum.mit.edu

