

Enhanced Flow Monitoring with P4 Generated Flexible Packet Parser

Tomas Cejka¹, Petr Velan¹, Jiri Havranek², and Pavel Benacek¹

¹ CESNET, a.l.e.

Zikova 4, 160 00 Prague 6, Czech Republic
{cejkat, petr.velan, benacek}@cesnet.cz

² CTU in Prague, FIT

Thakurova 9, 160 00 Prague 6, Czech Republic
{havraji6}@fit.cvut.cz

Abstract. Passive network flow monitoring provides visibility into network traffic. It is necessary for many applications such as accounting, network management, and security. As its origins are in packet switching and routing devices, the common flow exporter implementations process only necessary packet headers. Link layer protocols are often skipped, and only the first network and transport layer headers are used to construct flow records. However, the network traffic is gradually becoming much more complex as new protocols are being used in practice. We present a novel multi-layer flow monitoring approach that handles complex protocol encapsulation. To process packets with an arbitrary number of protocols, we have created a new packet parser based on the P4 language, which is easily extensible and widely used in SDN networks. We argue that the new multi-layer flow monitoring approach provides more precise and detailed statistics about the traffic of overlay networks at a backbone level.

1 Introduction

Traffic monitoring is necessary for many daily objectives such as accounting, operational awareness, or network security. Without enough network visibility, it is very difficult to detect network-wide malicious behavior and breach attempts. Flow monitoring has been used for this purpose for more than two decades, and significant research effort has been invested to attack detection based on flow data [24].

Flow records, which contain information about communication between two network devices, are exported by monitoring probes (flow exporters) and stored (and processed) by flow collectors. Each flow record is usually identified by an n-tuple consisting of timestamps, network addresses, transport protocol, and transport ports. The IP Flow Information eXport (IPFIX) standard [8] describes not only the protocol for flow export from the probes to the collectors but the flow monitoring architecture as well [23]. A straightforward approach, and also a commonly used one, is to get the flow record identifiers from the first visible

network and transport layers. The rest of packets is assumed to be a payload, which is aggregated to just a number of bytes.

In practice, low-level network infrastructure does not change fast enough as the new technologies evolve. As a result, we can see multiple network layer protocols encapsulated in a single packet. The reason is simple, the current infrastructure and logical network are used to carry a new overlay infrastructure. The aim of this paper is to provide more detailed information about such encapsulated traffic with the (almost) same flow processing tools so that it is possible to use existing analysis and detection methods.

Besides monitoring highly encapsulated traffic, we also deal with a lack of flexibility of the current flow exporters. The set of supported protocols and, in addition, the set of extracted header fields is fixed and hard-coded in the source codes of flow exporters. Usually, the functionality and the range of provided information might be extended by external plugins (linkable objects); however, we show an entirely different approach which uses the Programming Protocol-independent Packet Processors (P4) language to construct a highly flexible flow exporter. The presented flow exporter provides flexibility of P4 and also allows for deep packet inspection and traversing multiple encapsulated layers of the packets

This paper provides a novel approach to monitoring of a network traffic with multiple encapsulated protocols. Firstly, we present a generic parser based on the P4 language that is capable of easy extension of the supported protocols at the input of the flow exporter and extension of information contained in flow records. Secondly, we use this parser to construct a flow exporter that can process an arbitrary number of encapsulated layers. The new flow exporter creates a tree of flow records so that each encapsulated layer can be analyzed with minimal or no change to existing flow collectors and tools for analysis.

The rest of the paper is structured as follows. Sec. 2 presents related work in the area of flow exporters and P4 described network devices and explains the difference between the existing approaches and our approach. Sec. 3 describes usage of P4 for definition of the supported protocols, and construction of automatically generated flexible packet parser capable of recursive extraction of the encapsulated protocols. Sec. 4 proposes a multi-layer monitoring approach and describes its application in the new flow exporter and other existing flow-based tools. Sec. 5 discusses benefits and possible applications of the presented approach. Sec. 6 concludes the paper and lists our future work.

2 Related Work

Packet header parsing is the basic functionality of every network device. Nowadays, devices support lots of protocols; however, the purpose of the parsing is to use just information contained in the headers for switching, routing, or any other actions. This is the reason why the parsing usually ends up with the first needed occurrence of the protocol in the packet. This approach is the main principle of encapsulation and decapsulation used in the ISO/OSI model.

The first part of this section surveys existing work on flow monitoring of encapsulated traffic. The second part describes packet parsers and focuses especially on the P4 language.

2.1 Flow Exporters

An approach similar to the ISO/OSI principle is used for the current flow exporters, which usually utilize the first network and transport layers found in a packet. Additionally, it is possible to skip certain lowest layers such as VLAN, TRILL, or MPLS in order to analyze packets transferred through the network infrastructure. However, current implementations of flow monitoring tools do not support the representation of repeated protocols in a single packet.

A special category of devices and their parsers are L7 filters and analyzers (e.g., [9, 22, 27]). Such devices are capable of advanced traversing through the encapsulated header, however, they usually support only a limited number of combinations of the protocols in a packet. This can work in practice where the parsing and header traversal is partially done by the underlying layers of the L7 device. Automatic generation of protocol parsers is not a novel approach. In 2006, Pang et al. proposed a high-level language called binpac to describe parsers of L7 protocols in [21]. Contrary to binpac, this paper is focused on lower layers and their multiple occurrences in a packet.

Elich et al. in [11] investigate IPv6 tunneled traffic in a real network. The authors of the paper extend an existing flow exporter to examine and report values of specific fields of chosen IPv6 tunneling protocols. The additional fields are exported as a part of the original flow record that contains the tunneled traffic as a payload. The authors define custom IPFIX elements to represent the new fields and extend a flow collector to process these elements. Our work aims to provide support for general packet encapsulation, not only IPv6 tunnels. Moreover, we choose a different method of exporting information about encapsulated traffic.

The complexity of link layer protocols is steadily increasing. Traffic encapsulation is quite common whenever virtual networks are created or IPv6 transition mechanisms are applied. For instance, VLANs or similar mechanisms help to segregate a network into isolated parts due to many reasons such as security or routing. Most flow exporters provide the capability to process MPLS and VLAN tags, fewer provide mechanisms to unwrap tunneling mechanisms such as the GRE protocol. Although there are several open-source flow exporters, such as [6, 10, 15, 17] as well as many commercial ones [7, 12], we are currently not aware of any flow exporter that would allow processing of traffic with arbitrary encapsulation. Moreover, the structure of flow records is usually flat and adding dynamic fields for every encountered encapsulation layer is not supported by flow collectors either.

2.2 Packet Parsers

Software packet parsers and parsers created for operation in hardware switches and routers are different in their nature. Both can be optimized for speed; however, the hardware can take advantage of a parallel data processing which allows more efficient analysis of the packets. The hardware parsers usually need a parse graph of protocols that are supported and which shows the expected order of their presence in the packets. This limitation makes it difficult for packet parsers in networking devices to process arbitrary protocol encapsulation.

Gibb et al. studied design of packet parsers for hardware targets in [14]. The paper shows trade-offs in parser designs, identifies design principles for switches and routers, and describes a parser generator that outputs the source code in Verilog. The authors show that the parse graphs can be either fixed or programmable. However, even though the programmable parse graphs are easy to modify, they must be known at a compile time.

The authors of [16] present a parser design which allows processing of multiple headers in a single bound. They use a fast header field extractor with a lookahead function and implement a translation for NetFPGA devices. However, parser instructions are compiled from parse graphs as well, which limits the parser versatility.

There are many software packet parsing libraries for various programming languages. The ones which implement the parsers directly are not easily extensible and are not useful for our purpose. An example of such a library is libtins [13], which is a high-level, multiplatform C++ network packet sniffing and crafting library. However, some parsers can be generated from a high-level description or grammar, which allows new protocols to be added easily and efficiently. For example, Bangert and Zeldovich created a tool called Nail [2] for parsing and generating data formats. The main purpose of the tool is to generate safe parsers. By eliminating semantics actions, the authors were able to simplify the parsers and create generators as well. However, the expressive power of their grammar was reduced.

P4 [4] is an open source high-level language which was firstly introduced by Bosshart et al. in 2014. The language allows us to describe all aspects of a packet processing device (i.e. structure of the supported protocols, parsing packet headers and processing packets). The language is designed to be platform independent which makes it suitable for a general description of a network device. The P4 language is the next step in an evolving ecosystem of Software Defined Networks (SDN). The strongest idea of P4 is to provide an easy way for a rapid data plane prototyping, while the former implementations of SDN switches typically contained only a static set of supported actions and protocols. Even though P4 improves the flexibility of SDN switches, due to performance requirements, it is not usual to map the described parser of protocols with a limited number of supported layers.

The P4 language is the major domain-specific language for the description of network devices. The P4 language Consortium [19] provides several tools such as a P4 compiler frontend [20], a software tool for running translated programs [18],

and others. Current works show that a P4 description can be mapped on various technologies. For instance, Benacek provided a theory about the translation of P4 programs to FPGA target in [3], whereas the aim is to process high-speed (more than 100 Gb/s) network traffic. Open vSwitch [1], which is an open source software implementation of an OpenFlow switch, can be used as a software target of P4 programs.

3 Flexible Packet Parser Based on P4

A general architecture of the proposed software flow exporter is inspired by existing tools and our many-years experiences. The high-level architecture of the flow exporter is shown in Fig. 1 and can be divided into three main parts: *packet parser*, *flow cache*, and *flow export*. The main focus of this section lies in the *packet parser* block.

The flexibility of the flow exporter is mainly in the *packet parser* that is generated from a high-level description in P4. For our purpose, a subset of the language features is used to specify the list of supported protocols and possible transitions between them.




Fig. 1. The process of P4 program compilation and packet parser usage.

3.1 New Backend for P4-16 Compiler

P4 Language Consortium provides and maintains an open source compiler for the P4 language at the GitHub repository [20]. It must be noted, that there are two versions of P4 compiler, whereas the latest one (for the P4 version 16) is used in our work. Fig. 1 shows the compiler consisting of *frontend*, *midend*, and *backend*. The *frontend* is a parser of the supplied P4 source codes, which creates an internal tree representation (IR). The *midend* is responsible for optimizations

of IR driven by architecture dependent policies and, subsequently, IR is passed into the *backend*. The *backend* generates output for the target.

Originally, the P4 compiler contained some existing backend modules; however, for our purposes, we needed a different kind of output — source codes of a linkable parser that can be used in the flow exporter. Therefore, a new *backend*³ has been developed in order to automatically generate source codes of the *packet parser* in the C language⁴. As a result, generated *.c and *.h files contain a packet parser function, declaration of known protocol headers (**struct** in C), and useful constants. The source codes can be easily linked to any C project to handle packet parsing.

The generated parser is implemented as a finite state machine (inspired by existing eBPF backend) that expects the size of the packet and the array of its bytes as an input. A state is represented as a block of source code that consists of a header extraction and next state decision logic. The aim of the parser is to extract header fields from the packet. Since the description of transitions between different protocol headers might be complicated and dependent on bytes around the current position of a “packet cursor” (offset in the packet that is being processed), there are several function constructs in P4: **lookahead** (read bits without moving packet cursor), **advance** (move packet cursor without reading bits), and **extract**. These functions must be supported by the *backend* and are translated into the equivalent functionality of the *packet parser* in order to support complex header traversal.

3.2 Output of the Packet Parser

Besides flexibility of the set of supported protocols defined by P4, it is also possible to specify a list of protocol headers that can be returned as an output of the *packet parser*. Other protocol headers, which are defined in P4 but are not explicitly listed for output, can be parsed and skipped in a packet.

The output of the *packet parser* is shown in Fig. 2. Contrary to a traditional approach, which stops at the first recognized transport layer, the presented flexible parser returns all identified protocol headers as a linked-list. The linked-list output allows for a multi-layer approach of the flow exporter (more details are in the following sections).

4 Multi-Layer Flow Monitoring

We have presented a flexible packet parser based on P4 which allows us to process encapsulated network traffic, extract packet headers relevant to the flow monitoring, and skip headers that should be ignored. This section describes the implementation of a flow exporter based on this parser. We show how to extend

³ Publicly available at [5].

⁴ Former P4 compiler (version 14) was able to generate a *packet parser* in the form of some C source codes, however, functionality was limited compared to our new *packet parser*.




Fig. 2. Example of mapping the encapsulated multi-layer packet onto multiple flow records.

a flow cache to enable processing of an arbitrary number of encapsulation layers, how to aggregate the flows on these layers, and how to export the aggregated information using the IPFIX protocol. We also argue that the chosen approach requires relatively small or no change to existing flow collectors in order for them to process the flow records in the newly exported format.

New terminology is used throughout this section. An outermost header is the first header relevant to the flow monitoring that is found in the packet. An outermost flow is a flow identified by flow keys taken from the outermost packet headers. Note that although we talk about a single header that defines the flow, it is only a simplification as a flow keys are usually taken from both network and transport layer headers.

4.1 Flow Exporter Implementation

We have considered two options of representation of the encapsulated traffic in a flow cache — Extended flow records and Tree structure. Fig. 3 shows two packets with multiple IPv4 headers and both representations as well. The option of Extended flow record is to keep a single flow record based on the outermost headers (e.g., src IP 192.168.1.1 in Fig. 3) and represent the upper layers by added information fields. This can be easily implemented by defining flow keys to span over all layers that are present in the packet. The packets with the same values of flow keys belong to the same flow; therefore, the outer flow is split naturally. This approach was used by Elich et al. in [11]. The disadvantage of this approach is that each flow record has to be dynamically sized for it has to be able to contain information from an arbitrary number of layers. Moreover, the processing of such flow records on the flow collector would be cumbersome,




Fig. 3. Two options of representing encapsulated flows in a flow cache.

and the flow collectors would have to be significantly modified to support queries over multiple layers in the same flow record.

The second option, the preferred one, is to represent the multi-layer hierarchy of encapsulated flows as a tree with the outermost flow being the root of the tree. Each layer is represented by a standalone flow record that contains an identification of the parent flow. A flow record of the outermost layer has naturally undefined parent flow, i.e., it can be set to 0 or some other constant. This representation of flows has several advantages. The first is that each flow record can have the same structure with a fixed length. This makes the implementation of the flow cache easier since a simple hash table can be used. The only change to the flow cache is that each flow record must now contain an identifier of its parent flow or an indication that no such parent flow exists, i.e., the flow is the outermost one. This allows utilizing the existing flow cache implementations with minimal change.

We have decided to use the second option for the reasons listed above. Therefore, we need to find a unique identifier for each flow record so that it can be referenced from the child flows. A simple sequence of numbers could be used for this purpose; however, its state would have to be kept so that a restart of the flow exported does not reset the counter which would create duplicated identifiers. A simple solution to keep the state and solve collisions is to use a timestamp as a start of a sequence. Using a timestamp of the start of the flow as a key would work as well. However, the timestamp precision has to be high enough to ensure that no two packets could receive the same timestamp on a given medium (the timestamp must increase faster than the packets are received). Note that the

flow identifier collisions need to be handled by a collector as well since it can receive data from multiple exporters.

The current trend in flow monitoring is to gather as much information from the traffic as possible. The flexible P4 packet parser allows extracting values of any header field of the supported protocols. The values are then accessible in the linked list of the extracted headers and can be easily added to the flow record. For example, if we decide to analyze TCP window size, it is very easy to make the value accessible just by adding the attribute to the TCP header definition in the P4.

Unfortunately, P4 is not designed to process application layer headers which are often text-based or without a fixed structure. However, a pointer to the application layer is easily extracted by the P4, and the parsing process can be complemented by the use of any of the number of available application parsers.

4.2 Flow Collector Compatibility

When a flow collector receives a flow record with unknown elements, it usually discards only these elements and processes the rest of the flow. Therefore, multi-layer flow records can be easily processed by standard flow collectors. However, without the access to the *parent ID* element outermost flow records cannot be distinguished, which could have an undesirable impact on further processing, such as computing statistics. As a workaround and indication of the outermost flows can be exported as a supported yet otherwise unused field, that can be processed by the collector. However, this approach may lead to confusion, errors and is better to be avoided by using a more flexible flow collector.

To make a flow collector fully compatible with the multi-layer flow monitoring, it needs to be flexible to allow a definition of new elements. *flow ID* and *parent ID* elements must be properly received from the flow exporter, stored and must be accessible for later processing. An example of such a flexible collector is IPFIXcol [26] with FastBit storage backend and associated fbitdump query tool [25]. To avoid collisions during parent flow lookup, a flow source identifier must be used together with the parent flow identifier since there might be several flows with the same flow identifier created by different flow exporters.

5 Discussion

Since we have designed our flexible packet parser as a C library, it can be used for other purposes than flow monitoring. A network diagnostics systems might use the parser to automatically discover issues in encapsulated traffic. Host-based firewalls are another example of devices which can make use of the parser. Filter rules to match encapsulated packets are not usually supported, and they have to be specified using byte sequences as the encapsulated traffic is considered a payload by the firewalls. Our parser can enable firewalls to inspect all protocol headers and match even the encapsulated traffic easily.

We believe that a deployment of multi-layer flow monitoring in practice will bring new insights to the network traffic and help to identify hidden threats and issues. Since the required modifications to the flow collectors are minimal, it should not be difficult to deploy. However, in order to get multi-layer flow monitoring widely adopted, the existing IPFIX flow definition, which is widely used, might need to be updated.

The performance of our parser should be increased before the multi-layer flow monitoring is deployed. Our parser is not optimized for speed at this time; therefore it is not suitable for high-speed applications. However, it can still be used to monitor local networks and analyze packet captures. The overhead of storing more layers of flows in the flow cache can hinder performance since encapsulated traffic requires an update of multiple flow records for a single packet. The real impact of this overhead depends on the structure of the traffic and should be evaluated before the multi-layer flow monitoring is deployed in a high-speed network.

6 Conclusion

Flexibility is one of the key features of modern network devices. The P4 language, which was used in our work, was developed in 2014 in order to make the packet processing much more extensible and adaptable according to application needs. This paper is focused on applying benefits of the high-level P4 description in the domain of flow-based monitoring.

Sec. 2 describes the current state-of-the-art in the area of flow exporters, which are the base components of a flow-based monitoring infrastructure. Our research and many-years experiences from practice show that flow exporters, as well as many other network devices, do not parse the whole packets to look up all known protocol headers. This approach is natural and logical since any additional parsing costs processing power, which is limited and expensive in the network devices. On the other hand, analysis of the higher layers of the protocol headers increases visibility that is very important for various use-cases such as performance analysis or network security.

The main contributions of this paper lie in the usage of the P4 language and improvement of a flow exporter that applies a “multi-layer” monitoring approach. The P4 language is used for automatic creation of a packet parser with a flexible set of the supported protocols. This flexible packet parser is able to traverse through the all known protocol headers in a packet and pass the gathered information to the rest of the flow exporter.

Since the official compiler of the latest version of P4 (P4-16) does not support output format that we could easily use in our flow exporter, a new backend for this compiler was created, and it is available at our GitHub repository [5]. The backend generates C source codes that can be used in any project.

A novel approach of the “multi-layer” monitoring and its representation in a standard IPFIX format was explained in Sec. 4. It is based on the linked-list of identified protocol headers, which is the output of the flexible packet parser.

The goal is to represent different layers of encapsulation in a single packet using slightly extended IPFIX flow records. Despite the flow records being extended, existing tools might be still used for processing. In addition, this approach allows us to gain more information from a single packet compared to a standard IPFIX record.

As a future work, we plan to elaborate on the performance of the generated packet parser because the maximal throughput is a key feature that affects deployment of the flow exporter in high-speed networks. A list of supported protocols and extracted header fields can be defined in P4, however, the whole concept of partially generated flow exporter can be improved to support changes at runtime.

Acknowledgment

This work was supported by the project No. TH02010214 *Platform for Acceleration of Network Functions Virtualization* funded by the Technology Agency of the Czech Republic, internal grant No. SGS17/212/OHK3/3T/18 of CTU in Prague, and the project Reg. No. CZ.02.1.01/0.0/0.0/16.013/0001797 co-funded by the MEYS and ERDF.

References

1. A Linux Foundation Collaborative Project: Open vSwitch, <http://openvswitch.org/>
2. Bangert, J., Zeldovich, N.: Nail: A practical tool for parsing and generating data formats. In: Proceedings of the 11th USENIX Conference on Operating Systems Design and Implementation. pp. 615–628. OSDI’14, USENIX Association, Berkeley, CA, USA (2014), <http://dl.acm.org/citation.cfm?id=2685048.2685098>
3. Benacek, P.: Generation of High-Speed Network Device from High-Level Description. Ph.D. thesis, Czech Technical University in Prague (2017), <https://github.com/benycze/PhD-Thesis>
4. Bosshart, P., et al.: P4: Programming Protocol-independent Packet Processors. SIGCOMM Comput. Commun. Rev. pp. 87–95 (2014), <http://doi.acm.org/10.1145/2656877.2656890>
5. CESNET, a.l.e.: Flexible Packet Parser — Backend for p4c, <https://github.com/CESNET/fpp-p4c-backend>
6. CESNET, a.l.e.: flow_meter, https://github.com/CESNET/Nemea-Modules/tree/master/flow_meter
7. Cisco Systems, Inc., San Jose, CA and USA: Cisco Application Visibility and Control (AVC), <https://www.cisco.com/c/en/us/products/routers/avc-control.html>
8. Claise, B., Trammell, B.: Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information. RFC 7011 (Sep 2013), <https://rfc-editor.org/rfc/rfc7011.txt>
9. ClearFoundation: l7-filter (2017), <http://l7-filter.clearos.com/>
10. Deri, L.: nProbe: an Open Source NetFlow Probe for Gigabit Networks. In: Proceedings of the TERENA Networking Conference, TNC’03. Zagreb, Croatia (2003), <http://luca.ntop.org/nProbe.pdf>

11. Elich, M., et al.: An investigation into teredo and 6to4 transition mechanisms: Traffic analysis. In: 38th Annual IEEE Conference on Local Computer Networks - Workshops. pp. 1018–1024 (Oct 2013)
12. Flowmon Networks: Flowmon Probe, <https://www.flowmon.com/en/products/flowmon/probe>
13. Fontanini, M.: libtins, <http://libtins.github.io/>
14. Gibb, G., Varghese, G., Horowitz, M., McKeown, N.: Design principles for packet parsers. In: Architectures for Networking and Communications Systems. pp. 13–24 (Oct 2013)
15. Inacio, C.M., Trammell, B.: YAF: Yet Another Flowmeter. In: Proceedings of the 24th International Conference on Large Installation System Administration. pp. 1–16. LISA’10, USENIX Association, Berkeley, CA, USA (2010), <http://dl.acm.org/citation.cfm?id=1924976.1924987>
16. Kozanitis, C., et al.: Leaping multiple headers in a single bound: Wire-speed parsing using the kangaroo system. In: Proceedings of the 29th Conference on Information Communications. pp. 830–838. INFOCOM’10, IEEE Press, Piscataway, NJ, USA (2010), <http://dl.acm.org/citation.cfm?id=1833515.1833654>
17. Miller, D.: Softflowd, <https://www.mindrot.org/projects/softflowd/>
18. P4 Language Consortium: Behavioral Model Repository (Oct 2017), <https://github.com/p4lang/p4c>
19. P4 Language Consortium: P4 (2017), <https://p4.org/>
20. P4 Language Consortium: p4c (2017), <https://github.com/p4lang/p4c>
21. Pang, R., Paxson, V., Sommer, R., Peterson, L.: binpac: A yacc for writing application protocol parsers. In: Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement. pp. 289–300. ACM (2006)
22. Radware, Ltd.: Web Application Firewall (WAF) & Network Security Solution (2017), <https://www.radware.com/Products/AppWall/>
23. Sadasivan, G., et al.: Architecture for IP Flow Information Export. RFC 5470 (Informational) (Mar 2009), <https://www.rfc-editor.org/rfc/rfc5470.txt>, updated by RFC 6183
24. Umer, M.F., Sher, M., Bi, Y.: Flow-based intrusion detection: Techniques and challenges. Computers & Security 70, 238 – 254 (2017), <http://www.sciencedirect.com/science/article/pii/S0167404817301165>
25. Velan, P.: fbitdump, <https://github.com/CESNET/ipfixcol/tree/master/tools/fbitdump>
26. Velan, P., Krejčí, R.: Flow Information Storage Assessment Using IPFIXcol. In: Dependable Networks and Services. Lecture Notes in Computer Science, vol. 7279, pp. 155–158. Springer Berlin Heidelberg, Heidelberg (2012), http://dx.doi.org/10.1007/978-3-642-30633-4_21
27. Wireshark Foundation: Wireshark (2017), <https://www.wireshark.org/>