A Survey on Network Security Monitoring Implementations

Ibrahim Ghafir, Jakub Svoboda, Vaclav Prenosil

Abstract— **Network monitoring is a difficult and demanding task that is a vital part of a network administrator's job. Network administrators are constantly striving to maintain smooth operation of their networks. If a network were to be down even for a small period of time, productivity within a company would decline, and in the case of public service departments the ability to provide essential services would be compromised. There are different network security approaches. This paper provides the readers with an overview of concrete software implementations of the current network monitoring approaches. In addition, it presents a comparison between those implementations.**

Keywords—**Network security monitoring, packet capture, deep packet inspection, flow observation.**

I. Introduction

Monitoring helps network and systems administrators to identify possible issues before they affect business continuity and to find the root cause of problems when something goes wrong in the network. Whether it is a small business with less than 50 nodes or a large enterprise with more than 1000 nodes, continuous monitoring helps to develop and maintain a high performing network with little downtime.

For network monitoring to be a valuable addition to a network, the monitoring design should adopt basic principles. The monitoring system should be comprehensive and cover every aspect of an enterprise, such as the network and connectivity, systems as well as security. It would also be preferable if the system provides a single-pane-of-glass view into everything about the network and includes reporting, problem detection, resolution, and network maintenance. Further, every monitoring system should provide reports that can cater to a different level of audiences, the network and systems admin, as well as to management. Most importantly, a monitoring system should not be too complex to understand and use, nor should it lack basic reporting and drill down functionalities.

Network monitoring is a difficult and demanding task that is a vital part of a network administrator's job. Network administrators are constantly striving to maintain smooth

_____

Ibrahim Ghafir [1], Jakub Svoboda [2], Vaclav Prenosil [1]

[1] Faculty of Informatics, Masaryk University
[2] Institute of Computer Science, Masaryk University
Brno, Czech Republic
ibrahim_ghafir@hotmail.com, svob.jak@gmail.com, prenosil@fi.muni.cz

operation of their networks. If a network were to be down even for a small period of time, productivity within a company would decline, and in the case of public service departments the ability to provide essential services would be compromised. In order to be proactive rather than reactive, administrators need to monitor traffic movement and performance throughout the network and verify that security breaches do not occur within the network.

There are different network security approaches [1]. This paper provides the readers with an overview of concrete software implementations of the current network monitoring approaches. In addition, it presents a comparison between those implementations.

The remainder of this paper is organized as follows. Section 2 classifies the current network security monitoring implementations into three main classes. A comparison between presented implementations is provided in Section 3 and Section 4 concludes the paper.

II. Network Security Monitoring Implementations

This section classifies the current network security monitoring implementations into packet capture representatives, deep packet inspection representatives and flow-based observation representatives. It also provides information about the usefulness of particular tools for development of new network traffic analysis methods.

A. Packet Capture Representatives

1) Tcpdump

Tcpdump is a command line tool for packet capture analysis. Tcpdump can analyze both live traffic using the libpcap library and captured packet traces in PCAP format. Packets may be filtered both before and after the capture. Filtering before the capture can be done using BPF (Berkeley Packet Filter). Filtering after capture can be achieved using tcpdump's filters, described later in this section.

Data are printed out in text format. The output displays individual packets with information that include source and destination addresses, L4 protocol used, and L4 protocol flags. Figure 1 shows output listing two packets.

Packets to be displayed can be filtered using expressions. Filters can be imposed on source and destination addresses, ports, L3 and L4 protocols, and L4 protocol flags. Addresses

can be expressed in format of individual addresses or in CIDR notation. Multiple rules in the expression can be composed using boolean operators. The example on Figure 1 uses three filters. *src host 10.0.2.15* selects only packets originating in the IP address 10.0.2.15. *dst port 22* selects only packets destined for the port 22. Finally, *tcp[13] = 2* selects packets in which the decimal value of the 14th byte is 2. The filters are composed using the *and* word, which means only packets that meet all the criteria pass through the filter.

```
$ tcpdump -r pcap -n \
"src host 10.0.2.15 and dst port 22 and tcp[13] = 2"
reading from file pcapfile, link-type EN10MB (Ethernet)
20:20:40.512613 IP 10.0.2.15.54346 > 192.168.1.42.22: Flags [S],
seq 3123841387, win 29200, options [mss 1460,sackOK,TS val 509640
ecr 0,nop,wscale 7], length 0
20:20:41.356843 IP 10.0.2.15.45749 > 192.168.1.41.22: Flags [S],
seq 1764864395, win 29200, options [mss 1460,sackOK,TS val 509851
ecr 0,nop,wscale 7], length 0
```

Figure. 1. Example of a tcpdump filter and tcpdump's output.

Tcpdump needs root privileges to open the network interface. Operation without full root is possible using SUID or Linux capabilities. Granting the *tcpdump* executable *cap_net_raw* and *cap_net_admin* capabilities allows tcpdump to be run as a regular user.

2) Wireshark

Wireshark is a graphical tool for packet capture analysis. While Wireshark and tcpdump are implementations of the same architectural approach, their underlying ideas differ. Tcpdump is as close to the raw data as possible, while Wireshark strives to provide higher-level representation of the same data.

Wireshark can analyze both live traffic using the libpcap library and captured packet traces in PCAP format. Captured packets can be filtered both during and after the capture. Filtering after capture can be achieved using filter expressions. Filtering during capture can be done using BPF.

Data are displayed as text arranged in a scrollable colored list and expandable boxes. Wireshark's main window has two frames. The upper frame displays list of captured packets with their basic attributes displayed. A line may be colored, based on the protocol the individual packet belongs to. When the user selects packet in the upper frame, this particular packet is displayed in the lower frame. The lower frame's representation includes several boxes that can be expanded and collapsed. The boxes contain various attributes of the packet as well as representations of the packet's data in ISO OSI layers' data the packet is part of. Also, related data from other packets may be displayed there, HTTP TCP stream for instance.

Packets displayed in the upper frame can be filtered using expressions entered in the text box on the top of the window. The expression vocabulary is richer that the one of tcpdump. Figure 2 shows the architecture of Wireshark.

Wireshark uses a separate program to capture the traffic, dumpcap. The reason is to allow separation of privileges [2]. Wireshark can be run as a regular user and only dumpcap has to be given special permissions. Dumpcap needs root privileges to open the network interface. Operation without the user having root access is possible using either SUID or Linux capabilities. Granting the *dumpcap* executable *cap_net_raw* and *cap_net_admin* capabilities allows dumpcap to be run as a regular user without SUID.


Figure. 2. Wireshark architecture showing privilege separation.

B. *Deep Packet Inspection Representatives*

1) Snort

Snort is an intrusion detection system performing deep packet inspection using pattern matching. The pattern matching is implemented in the form of rules [3]. Rules are structured text files describing network traffic data of interest. Typically, rules are used to generate alerts when a security-related incident occurs, such as malware activity, attack, or breach of security policies. A rule contains information specifying when the rule should be triggered. An important part of these information is one or more patterns that are searched in the network traffic. The pattern can be a sequence of characters or bytes or a regular expression. Figure 3 shows Snort rule structure.

rule header	alert tcp any any -> 192.168.1.0/24 111 \
rule options	(content:"|00 01 86 a5|"; msg:"mountd access";)

Figure. 3. Snort rule structure

Snort rules have a specific structure. The beginning of the rule before the parentheses describes which network flows the rule refers to. This is called the rule header. The rule header specifies the action the rule should perform (*alert* for instance), L3 protocol and source/destination IP addresses and ports on which to match. Variables may be used in place of IP addresses. The rest of the rule inside the parentheses is called the rule options. Rule options specify content on which the rule matches and other properties of the rule, its name, classification type, etc. The most important part of the rule is the *content* keyword that specifies a pattern to be found in the packet's payload. The *content* keyword's function can be further changed using modifiers. For instance, modifiers offset, distance, depth, and *within* control in which areas of the packet the rules are matched [4].

Snort has three special keywords, *byte_test, byte_jump,* and *byte_extract*, that allow to adjust the pattern matching based on data in an individual packet [5]. The first two keywords behave as patterns that match when their conditions are true. *byte_test* performs arithmetic ($<$, $\leq$, $=$, $>$, $\geq$) and bitwise (AND, OR) comparisons on sequences of bytes. *byte_jump* puts a space before the next pattern with size inferred from payload's byte value. If the jump is possible, *byte_jump* also behaves as a match. This behavior can be used to match packets with a specific length based on specific data in the payload. *byte_extract* converts specified bytes into a numerical variable that can be used later in the rule. These keywords do not allow more complicated decoding or processing.

The Snort's architecture allows the implementation of so-called preprocessors [6]. Preprocessors read the packet before rule evaluation, serially in the order specified by Snort's configuration. This allows implementation of additional rule keywords. Moreover, preprocessors allow implementation of functionality more complicated than just pattern matching, such as data decoding and anomaly detection. For instance, the Normalizer preprocessor converts equivalent values to a unified format with the goal of making IDS evasion harder. Snort's architecture including the preprocessors is depicted in Figure 4.
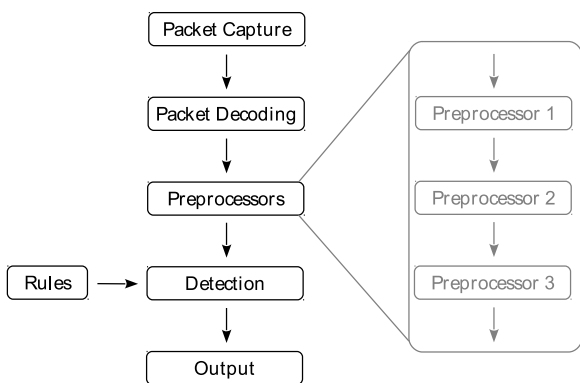

Figure. 4. Snort architecture.

There are several preprocessors for anomaly detection available. Frag3 and stream5 preprocessors are integrated in the official Snort distribution and detect protocol anomalies. SPADE [7], PHAD [8], and snortad [9] are 3rd party preprocessors detecting traffic anomalies. Snort preprocessors are usually implemented in C. They allow implementation of similar concepts to those that can be implemented in the event-based architecture

Snort needs root privileges to open the network interface. It is possible to configure Snort to drop its privileges to a non-root user once it opens the network interface.

Snort is a single-threaded application. Multithreaded Snort setups work in the following way: The monitored traffic is divided by flows into multiple parts; each part of the traffic is fed to a single Snort instance [10].

2) Suricata

Suricata is an intrusion detection system performing deep packet inspection using pattern matching. Figure 5 shows Suricata rule structure, the rule header and rule options are as same as the Snort rule structure.

rule header   alert tcp any any -> 192.168.1.0/24 111 \
rule options   (content:"|00 01 86 a5|"; msg:"mountd access";)

Figure. 5. Suricata rule structure.

Suricata uses similar rules to Snort and is compatible with Snort rules. The rule structure is the same for both Snort and Suricata. The difference between the two is in the keywords and protocols that can be specified. Suricata allows specification of several L7 protocols on top of the L3 protocols supported by Snort, *http, ftp, tls, smb* and *dns* [11]. Some keywords behave differently than in Snort, for instance, the *fast_pattern:only* keyword doesn't make a difference in processing, unlike in Snort. Some keywords are supported only by Suricata, such as the *iprep* keyword for matching IP reputation data and the *dns_query* keyword for analyzing only the DNS response body.

Suricata's architecture is similar to Snort's one with a difference. What corresponds to the preprocessor part in the Snort's architecture is divided in two in Suricata, decoding and detection. We found out about this by studying the source code [12]. Decoding modules add information to the internal representation of packets in Suricata. Detection modules rely on this internal representation and provide keywords for use in rules. Overview of the Suricata's architecture is shown in Figure 6.


Figure. 6. Suricata's architecture.

Each packet is first processed in decoding functions and then in detection modules. Decoding functions read the packet and save the decoded data into an internal representation of the packet. The decoding functions are called one at a time on the packet. Extending the decoding functionality is possible by implementing a new decoding function and placing it into the decoding pipeline. The decoding pipeline starts with the source of captured packets, then L2 is decoded, and then protocols on higher layers are decoded.

Upon decoding, the packets pass detection. The detection is governed by rules and depends on the decoding step. The rules are matched with the internal packet representation. The matching process is broken into several detection modules in all of which the matching takes place. Unlike decoding, detection is parallelized and one packet can be processed in multiple detection modules at the same time. Extending the detection functionality is possible by implementing a new detection module and registering it in the table of detection methods.

Suricata is written in C and the modules for Suricata have to be written in C. There are no plans supporting C++. C requires greater programming expertise than the Bro language. Therefore, this property makes Suricata not the best prototyping tool available.

Suricata is multithreaded out of the box. Even though it is not as fast as Snort on a single-CPU computer, Suricata is designed to scale on computers with tens of CPUs [13]. The multithreading approach is different from Snort. Multithreaded Snort setups divide the monitored traffic by flows into multiple parts, each processed by an individual Snort instance. Suricata, on the other hand, does not require the traffic balancing since it manages multithreading itself. This approach makes it more user-friendly.

3) Bro

Bro [14] is a network security monitor performing deep packet inspection using event-based analysis. In contrast to Snort and Suricata, Bro is primarily not rule-driven. Instead, it implements a Turing-complete scripting environment [15]. Rule-based detection as well as arbitrary detection algorithms can be implemented in this environment. Bro detection rules are described by scripts. Figure 7 shows Bro architecture.
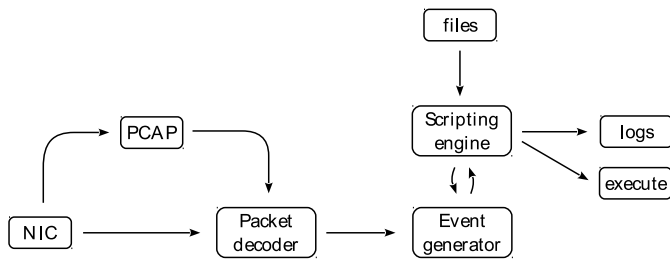

Figure. 7. Bro architecture.

The Bro's scripting environment uses the Bro programming language. It is an interpreted, typed language. What makes it special are domain-specific types. For example, the *addr* type holds an IP address [16]. Variables of structured types are reference type variables. This makes processing of large sets or tables efficient, since only the references are copied, not the data itself. There are two types of collections, sets and tables. Loops are available in the form of iteration through collections. The Bro programming language lacks other forms of loop control, presumably serving as a deterrent against overly complex algorithms. This is a reasonable requirement for network traffic monitoring when the

processing is done in real time. And that exactly is the most significant goal of Bro, to allow real-time network traffic analysis and save already processed results to log files.

The default installation of Bro contains many scripts implementing various sorts of traffic analysis. Some of the items the default Bro setup monitors are: Bidirectional flows, DHCP leases, DNS queries and responses, MD5 and SHA1 hashes of files transmitted over unencrypted protocols, HTTP requests and user agents, port scans, email headers from SMTP traffic, successful and unsuccessful SSH connections, SSL certificates, SYSLOG messages, traffic tunnels.

Since the preinstalled scripts usually expose an API in the form of events, they can be used by user scripts, extending the default functionality.

The core of Bro, implemented in C, processes network traffic, performs DPI and generates events about what is happening in the traffic. Events generated by the core are listed in the bif files [17]. Many events are generated, spanning L2 through L7. Examples are a new ARP packet, closed TCP connection, HTTP request, etc. In other words, this type of DPI performs semantic matching of network events instead of simple pattern matching, as opposed to Snort and Suricata. Majority of the events provide context, typically in the form of information about the relevant connection. The events are then processed by the Bro scripts.

Bro scripts use so-called event handlers to listen to the events. The usual reactions to events vary. On the one hand, the simplest possible processing saves the event information to a log file. On the other hand, some scripts implement fairly complex processing and generate additional types of events. This further extends DPI abilities of Bro. Scripts can handle events generated both by the core and by other scripts. Figure 8 shows a very short module that just writes "Hello world!" to the standard output when Bro starts.

```
module helloworld;
    event bro_init() {
        print "Hello world!";
    }
```

Figure. 8. A simple Hello world! script.

The scripting engine hosts the scripts and dispatches events generated both by the scripts and the core to the scripts listening to these events. It also allows operations like file access and execution of applications native to the operating system. This functionality can be used by advanced scripts. File access may be used to fetch information from external sources, e.g., a blacklist. Execution facility may be used for many purposes. One example is reporting issues to a ticket managing software via email. The *sendmail* executable can be used by such a script. Another example is automatic triggering of a remotely triggered black hole by executing a program that does the blackholing.

Bro scripts are organized in so-called modules. A module can be implemented wholly in one file or can be broken into several files. Two identifiers with the same name in two

different modules do not collide with each other. Cross-module references can be made using the name *name_of_module::name_of_identifier*.

A module can define types, variables, functions, and event handlers. These entities can be either local to the module or globally accessible from other modules.

It is possible to define custom types using *enum, set, table, vector,* and *record*. *Enum* in Bro is similar to *enum* in other languages. *Set* is similar to *HashSet<T>* in C# in its functionality [18], albeit the syntax is different. *Table* is similar to *Dictionary<TKey,TValue>* in C# [29] with the difference that C# allows only one key while Bro allows multiple keys. *Vector* is a table indexed by count. *Count* is the name for int in Bro. *Record* is similar to C# class [20] that contains only fields [21]. Both Bro *record* and C# *class* are reference types, meaning assignment of its instance copies only the reference (pointer), not the whole instance. This can be compared to C# *struct* which is a value type, meaning assignment of its instance copies the whole instance.

Bro can be run both as a single-threaded application and as a multithreaded distributed application. The single-threaded mode is called standalone while the multithreaded one is called cluster. If Bro is used as a platform for development of proof-of-concept methods, the standalone mode is usually more appropriate than the cluster mode. Development for the cluster mode is more difficult than for the standalone mode because additional functionality has to be used by scripts [22].

C. *Flow-based Observation Representatives*

Flow-based observation architecture contains two main components; a flow exporter and a flow collector. This section covers representative implementations of both flow exporters and flow collectors.

1) Flow Exporters

nProbe [23] is a commercial open-source flow exporter. Data can be exported in NetFlow v5, NetFlow v9, and IPFIX formats. nProbe has an application visibility (nDPI) ability, which is used for detection of application-specific protocols. This information is saved in a custom column in NetFlow v9 or IPFIX format. It is difficult to obtain nProbe source code for free.

YAF [24] is an open-source flow exporter. Data are exported in the IPFIX format. A passive OS fingerprinting functionality based on the p0f software can be compiled into YAF. YAF supports modules that implement DPI. However, YAF does not provide DPI in default setup.

QoF [25] is a fork of YAF. It removes all payload inspection abilities and instead focuses on passive performance measurements.

ipt-netflow [26] is a plugin for iptables for flow export. Data can be exported in NetFlow v5, NetFlow v9, and IPFIX formats. There is no special functionality besides standard network flows. There is also no apparent focus towards high-throughput networks. ipt-netflow is open-source.

pmacct [27] is an open-source flow exporter and flow collector. Data can be exported in NetFlow v5, NetFlow v9, sFlow v5, and IPFIX formats. Supports high-throughput networks using PF_RING. No DPI-related functionality is available in pmacct.

softflowd [28] is an open-source flow exporter performing export to NetFlow v1, v5, and v9 formats. There is no apparent effort to provide anything on top of regular NetFlow data export.

2) Flow Collectors

nProbe is not only a flow exporter, it is also a flow collector. Available storage backends are MySQL, SQLite, text files, and binary files. The nProbe flow collector was created because its author deemed other collectors available at the time to be too cumbersome to use.

IPFIXcol [29] is an IPFIX collector designed for high-throughput networks. IPFIXcol claims to be flexible. Storage backend can be customized using output plugins. IPFIXcol also allows implementation of so-called IPFIX mediators, used for processing of the collected data before it hits the collector.

flowd [30] is a NetFlow v1, v5, v7, and v9 collector. It is created under the UNIX philosophy to do just one thing. The collected data are saved in a binary format. flowd is provided with Perl and Python interfaces for reading the binary data. flowd strives for security using privilege separation of components. flowd is open-source and freely available.

nfdump [31] consists of several tools. The nfcapd tool listens to NetFlow v5, v7, v9 streams and saves them to nfcap files. The nfdump tool can be used for analysis of nfcap files. nfdump uses similar filter syntax to tcpdump. nfdump is open-source and freely available.

pmacct [32] as a collector has several storage backends available. It can use MySQL, PostgreSQL, SQLite, MongoDB, BerkeleyDB, and flat files. Among other formats, it can collect NetFlow v1-v9 and IPFIX.

SiLK [33] is a collector for NetFlow v5, v9, and IPFIX data. It is designed for high-throughput networks. SiLK consists of multiple tools and plugins for filtering, analysis, and processing of flow data.

III. Comparison

With respect to the selection of network traffic monitor suitable for DPI, the following criteria have been evaluated for each mentioned traffic monitor:

- *Prototyping:* Is the network traffic monitor suitable for creation of method prototypes?

- *Developer-friendliness:* Does the network traffic monitor allow development of new traffic analysis methods in an easy to use way?

- *Extensibility:* Is it possible to extend the existing functionality of the network traffic monitor in a reasonable way? What programming language does the API use?

The descriptions of individual network traffic monitors in this paper indicate answers to these criteria. Table I shows the summary.

TABLE I. BRO IS THE SUITABLE TOOL FOR CREATION OF PROTOTYPES

Monitor	Prototyping	Developer friendliness	Extensibility
Tcpdump	No	No	No API
Wireshark	No	No	No API
Snort	No	No	C language
Suricata	No	No	C language
Bro	Yes	Yes	Bro language

IV. **Conclusion**

There are several different approaches to network monitoring. Each approach is the best fit for a different purpose. Wireshark is good for manual analysis, predominantly of small capture files. Tcpdump is packet-oriented and works well in those use cases where filtering individual packets by L3/L4 attributes like IP address, TCP flags, payload bytes, etc. is sufficient. It does not work well for stream reassembly or L7 protocol analysis. Snort and Suricata work well when the goal is to match patterns in network data. Bro allows development of advanced detection methods.

Bro is the best software/environment for development of novel detection or processing techniques. It can be used for continuous monitoring of high-throughput networks. The scripting environment is extensible in a memory-safe language specialized in network data processing. It is not constrained by belonging to a single paradigm for network monitoring like the other tools. Unfamiliarity is a disadvantage, compared to more known tools like wireshark, tshark, snort and suricata.

References

[1] J. Svoboda, I. Ghafir, V. Prenosil. Network Monitoring Approaches: An Overview. In Proceedings of International Conference on Advances in Computing, Communication and Information Technology. Birmingham, UK, 2015. ISBN: 978-1-63248-061-3.

[2] Keuter, J.: Privilege Separation [online]. [cit. 2015-06-18]. URL, http://wiki.wireshark.org/Development/PrivilegeSeparation.

[3] Snort Syntax and Simple RuleWriting [online]. [cit. 2015-06-18]. URL, http://www.anotherchancecomputers.com/uncategorized/snort-syntax-and-simple-rule-writing/.

[4] Esler, J.: Offset, Depth, Distance, and Within [online]. URL, http://blog.joelesler.net/2010/03/offset-depthdistance-and-within.html.

[5] Writing Good Rules [online]. [cit. 2015-06-18]. URL, http://manual.snort.org/node36.html#testing_numerical_values.

[6] Esler, J.: Preprocessors [online]. [cit. 2015-06-18]. URL, http://manual.snort.org/node59.html.

[7] Biles, S.: Detecting the Unknown with Snort and the Statistical Packet Anomaly Detection Engine (SPADE) [online]. [cit. 2015-06-18]. URL, http://webpages.cs.luc.edu/~pld/courses/447/sum08/class6/biles.spade.pdf.

[8] Mahoney, M.: Network Anomaly Intrusion Detection Research at Florida Tech. [online]. URL, http://cs.fit.edu/~mmahoney/dist/.

[9] AnomalyDetection: Home - Snort.AD [online]. [cit. 2015-06-18]. URL, http://www.anomalydetection.info/?home,1.

[10] Verplanke, E.: Understand packet-processing performance when employing multicore processors [online]. [cit. 2015-06-18]. URL, http://www.embedded.com/design/connectivity/4007065/Understand-packet-processing-performance-whenemploying-multicore-processors.

[11] Suricata Rules [online]. [cit. 2015-06-18]. URL, https://redmine.openinfosecfoundation.org/projects/suricata/wiki/Suricata_Rules.

[12] suricata/src at master [online]. [cit. 2015-06-18]. URL, https://github.com/inliniac/suricata/tree/master/src.

[13] Julien, V.: On Suricata performance [online]. [cit. 2015-06-18]. URL, http://blog.inliniac.net/2010/07/22/on-suricataperformance/.

[14] Paxson, V.: Bro: A System for Detecting Network Intruders in Real-time. Comput. Netw., 1999: s. 2435 − 2463, ISSN 1389-1286, doi:10.1016/S1389-1286(99)00112-7.

[15] The Bro Network Security Monitor [online]. [cit. 2015-06-18]. URL, http://www.bro.org/documentation/overview.html.

[16] Types and attributes − Bro 2.2 documentation [online]. [cit. 2015-06-18]. URL, http://bro.org/sphinx/scripts/builtins.html.

[17] All Bro Scripts [online]. [cit. 2015-06-18]. URL, http://bro.icir.org/sphinx/scripts/scripts.html.

[18] Microsoft: HashSet<T> Class [online]. [cit. 2015-06-18]. URL, http://msdn.microsoft.com/en-us/library/bb359438.aspx.

[19] Microsoft: Dictionary<TKey, TValue> Class [online]. [cit. 2015-06-18]. URL, http://msdn.microsoft.com/en-us/library/xfhwa508\%28v=vs.110\%29.aspx.

[20] Microsoft: Classes and Structs (C# Programming Guide) [online]. [cit. 2015-06-18]. URL, http://msdn.microsoft.com/en-us/library/ms173109.aspx.

[21] Microsoft: Fields (C# Programming Guide) [online]. [cit. 2015-06-18]. URL, http://msdn.microsoft.com/en-us/library/ms173118.aspx.

[22] base/frameworks/cluster/main.bro [online]. [cit. 2015-06-18]. URL, https://www.bro.org/sphinx/scripts/base/frameworks/cluster/main.html.

[23] Deri, L.: nProbe: an Open Source NetFlow probe for Gigabit Networks. In In Proc. of Terena TNC 2003, 2003. URL, http://luca.ntop.org/nProbe.pdf.

[24] Christopher Inacio, B. T.: YAF: Yet Another Flowmeter [online]. [cit. 2015-06-18]. URL, http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.368.3172&rep=rep1&type=pdf.

[25] Trammell, B.: YAF-derived flow meter for passive performance measurement [online]. [cit. 2015-06-18]. URL, https://github.com/britram/qof.

[26] NetFlow iptables module [online]. [cit. 2015-06-18]. URL, http://sourceforge.net/projects/ipt-netflow/.

[27] Lucente, P.: pmacct: steps forward interface counters [online]. [cit. 2015-06-18]. URL, http://www.pmacct.net/pmacct-stepsforward.pdf.

[28] softflowd - A software NetFlow probe [online]. [cit. 2015-06-18]. URL, https://code.google.com/p/softflowd/.

[29] Velan, P.; Krejci, R.: Flow Information Storage Assessment Using IPFIXcol. In Dependable Networks and Services, editace R. Sadre; J. Novotny; P. Celeda; M.Waldburger; B. Stiller, Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2012. ISBN 978-3-642-30632-7, s. 155−158, doi:10.1007/978-3-642-30633-4_21.

[30] flowd - small, fast and secure NetFlow collector [online]. [cit. 2015-06-18]. URL, http://code.google.com/p/flowd/.

[31] NFDUMP [online]. [cit. 2015-06-18]. URL, http://nfdump.sourceforge.net/.

[32] Lucente, P.: pmacct project: IP accounting iconoclasm [online]. [cit. 2015-06-18]. URL, http://www.pmacct.net/.

[33] Gates, C.; Collins, M.; Duggan, M.; aj.: More Netflow Tools for Performance and Security. In Proceedings of the 18th USENIX Conference on System Administration, LISA' 04, Berkeley, CA, USA: USENIX Association, 2004, s. 121−132.