A Scalable Architecture for Network Traffic Forensics

Viliam Letavay
Faculty of Information Technology
Brno University of Technology
Brno 61266, CZ
Email: iletavay@fit.vutbr.cz

Jan Pluskal
Faculty of Information Technology
Brno University of Technology
Brno 61266, CZ
Email: ipluskal@fit.vutbr.cz

Ondřej Ryšavý
Faculty of Information Technology
Brno University of Technology
Brno 61266, CZ
Email: rysavy@fit.vutbr.cz

***Abstract*—Availability of high-speed Internet enables new opportunities for various cybercrime activities. Security administrators and LEA (Law Enforcement Agency) officers call for powerful tools capable of providing network communication analysis of an enormous amount of network traffic moreover, capable of analyzing an incomplete network data. Big data technologies were considered to implement tools for capturing, processing and storing packet traces representing network communication. Often, these systems are resource intensive requiring a significant amount of memory, computing power, and disk space. Presented paper describes a novel approach to real-time network traffic processing implemented in a distributed environment. The key difference to most existing systems is that the system is based on a lightweight actor model. The whole processing pipeline is represented in terms of actor nodes that can run in parallel. Also, actor-model offers a solution that is highly configurable and scalable. The preliminary evaluation of a prototype implementation supports these general statements.**

***Keywords*–Network forensic analysis; Network traffic processing; Actor model**

I. Introduction

The expansion of computer networks and Internet availability opens new opportunities for cybercrime activities moreover, security incidents associated with network applications. The amount of connected devices grows, and traffic speed increases. Security administrators and LEA officers call for powerful tools that enable them to extract useful information from network communication [1]. The network forensics that is responsible for capturing, collecting and network data analyzing is getting more important [2].

In the forensic investigation, the network traffic is continuously captured from multiple sources. The captured network data has a form of packet traces that have to be processed and analyzed up to the application layer. The network forensic tool has to decode protocols at different network layers of the TCP/IP model and various encapsulations. For LEA, interesting information lies in application messages such as instant messaging, emails, voice, RTP, localizable information, documents, pictures, etc. The form and relevance of extracted artifacts may differ from case to case. Often, communication is encrypted. In this case, meta-data can be the only a piece of information available. In all cases, the network forensic processing system has to be able to extract artifacts from the network traffic reliably, even if the packet capture is corrupted, for instance, some connections are incomplete, packets are malformed, or chunks of packets were not recorder because of capturing device issues.

The amount of data that needs to be processed to extract evidence from the network communication depends on the kind of a case that is investigated but usually gets large. It is very difficult to decode, extract and store the immense mass of information for further processing. We propose a distributed network forensic framework based on the *actor model* that is computation effective, and capable of linear scalability. Scalable properties of *actor model* design for network forensics are promising as shown by the VAST platform [3]. Similarly to VAST, our solution provides real-time data ingestion and interactive data analysis, but in addition to VAST, we consider the full artifact extraction up to the application layer. Although if it requires more computation resources, we demonstrate that if can still be achieved in a more straightforward and less resource consuming environment comparing to Apache Hadoop technology, which is the norm for big data processing.

In Section II, we describe tools used by network forensics practitioners. Section III addresses issues faced by investigators and our proposed solution, which architecture is broadly discussed in Section IV. Section V evaluates preliminary performance results, and Section VI concludes the paper.

II. Background & Related Work

Network forensics is a process that identifies, captures and analyzes network traffic. Network forensic techniques are used by several network forensic frameworks [4], [5], [6], [7], [8], [9] and tools intended for intrusion detection (Bro, Vast, Moloch) [10], [11], [12] , network security monitoring (Microsoft Network Monitor, TShark, Wireshark, tcpdump), and network forensic investigation for LEAs (Netfox Detective, PyFlag, NetworkMiner, EnCase, XPlico). Commonly available forensics tools are implemented either as a classic desktop or command line application or a traditional client-server.

To overcome the limitations of traditional tools, we propose to use distributed computing. The models for distributed processing [13], [14] are more suitable for real-time network forensic analysis from multiple sources, such as logs and captured communication. The models are based on an agent system, where numerous agents perform the collection task. The extracted information is sent to the forensic network server and analyzed on this single node [15] only. The *forensic server* is the bottleneck that has to process all the data. To avoid this bottleneck, the GRR live forensic system [16] utilizes a cluster of servers. The system deploys agents running on users' computers that provide access to forensic information, e.g., remote raw disk and memory access. Processing of forensic data is done as flows. Each flow is maintained on the server. Server nodes run workers that process the active flows. Adding more server nodes enable to run more workers and thus if it is possible to handle more clients simultaneously.

Elimination of bottlenecks in the architecture offers scalability and improved reliability. The *actor model* is one of the attractive solutions that address the problem elegantly and efficiently. It comes with a separate unit called an *actor*. Actors execute independently and in parallel. They communicate asynchronously via message passing, and their state is otherwise immutable. Actor's behavior determines how to process the incoming message. Actor system is the key enabler for the VAST system [3]. In VAST, actors implement importing, archiving, indexing and exporting processed data. Actors live in nodes that map to system processes. The system scales by creating more nodes either on the single machine or a cluster of computers.

Moloch is another tool, worth to mention, that uses principles of distributed computing for massive scale network traffic monitoring, full packet capturing and indexing. Moloch system consists of sensors that capture the communication and Elasticsearch database that is a distributed search and analytics engine. The system scales by adding new nodes running Elasticsearch instances.

III. Problem Statement and Solution

Our goal is to design and create a system capable of long-term, high-speed, real-time network traffic filtering and processing up to the application layer. The software solution should be scalable and hardware independent. To achieve this, we have to deal with the challenges elaborated in the rest of this section.

A. Architectural Design

How to create a system for packet filtering and analysis of communication that can identify application protocols, gets forensics artifacts and searches through them?

Network forensics is a tedious work that strictly relies on completeness and precision of all undertaken steps to gain a piece of a puzzle that fits together as a shred of evidence. Considering the current speeds of regular users' home network connection(s), a comprehensive classical analysis on a single machine would require enormous computation resources. Try to imagine, that each network packet would be analyzed by many protocol dissectors with a goal to extract, for example, an acknowledgment of email delivery. To achieve this goal, with optimal computational resources, we must revisit currently utilized methods and redesign them to work in a distributed environment which brings new challenges to architecture design, application of algorithms, data synchronization, and so on.

B. Scalability on Commodity Hardware

How can the solution be scalable and hardware independent despite the hardware limitations?

Let us consider this imaginary demonstration. The math is simple, one computer with $1\,\mathrm{Gbps}$ NIC that has a relatively simple task to capture traffic during full line load would be required to write to a disk under the constant speed of $1000\mathrm{Mbps} \approx 125\,\mathrm{MB/s}$. Our system has to guarantee that no data loss occurs during the capture. A suspect can simultaneously download and upload data which means that the monitoring device cannot have only one $1 * 1\,\mathrm{Gbps}$ NIC, but if needs $2{*}1\,\mathrm{Gbps}$ cards, one for uplink, one for downlink. Thus, the required speed of continuous disk writing would be $2 * 125\,\mathrm{MB/s} \approx 250\,\mathrm{MB/s}$. Now, if the requirement is to store the communication for one day, the disk capacity has to be $250\,\mathrm{MB/s} * 86\,400\,\mathrm{s} \approx 21.6\,\mathrm{TB}$. This is achievable with commodity hardware, e.g., $2 * 12\,\mathrm{TB}$ drives with RAID 0 or $4 * 12\,\mathrm{TB}$ with RAID 1+0 — assuming higher write/read speed than $250\,\mathrm{MB/s}$. However, what if only one day is not enough? For a typical forensic case, capturing period spawns through weeks or months.

From our previous experiments, we know that a single computation node is limited and commodity hardware is hardly sufficient to perform all required operations in real-time and over long periods. Separation of frames into a conversation which needs a dissection of the network protocols up to the application layer, which speed is roughly $300\,\mathrm{Mbps}$ [17, pp. 45-51] is not sufficient. On the other hand, we are confident that the application created and optimized for this singular purpose can do the processing faster and breach the $1\,\mathrm{Gbps}$ line speed. Nevertheless, we do not believe that a single machine solution with commodity hardware is capable of doing overall analysis and extraction of information from the application layer. We have to design our solution as a distributed system across multiple machines.

C. Overall Performance

What scalability and acceleration of data processing can be achieved?

The proposed solution is based on the actor model. Each actor represents an independent processing unit. The communication between actors is managed by messaging. The actor has no shared state; thus all actors can work in parallel. If actors run on the same node, the message passing has a little overhead compared to a function call or a loop. However, if actors scale over multiple nodes, messages need to be serialized. This process introduces latency and consumes part of processing power. The scalability of the actor model is linear [3].

IV. Architectural Design

Incomplete data provided by unreliable traffic interception can lead to inaccurate results; some information may be lost, some fabricated by reconstruction process [18]. Keeping the above facts in mind, the processing cannot strictly follow RFCs and behave like a *kernel* network stack implementation, but it has to incorporate several heuristics. For example, to fill missing gaps in data, and to consider these fillings during application protocol processing, or never to join multiple frames into a single conversation unless it passes more advanced heuristic-based checks. Network forensic tools that we have worked with do mostly respect RFCs and thus may produce misleading results as shows by Matousek et al. [18].

We propose a distributed architecture composed of commodity hardware that will be capable of linear scalability, and capable of efficient resource utilization. The overall architecture is shown in Figure 1.

At the top level, we have divided the entire process into the two main stages:

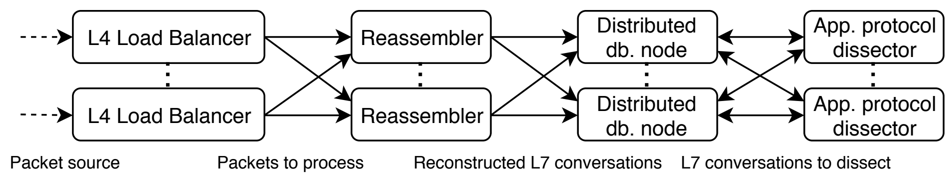

FIGURE 1. ARCHITECTURE DIAGRAM SHOWING THE PROPOSED SYSTEM NODES WITH INFORMATION FLOW BETWEEN THEM.

- Data preprocessing — The reconstruction of conversations at the application layer (L7), i.e., consecutive segregation of captured communication into internet (L3), transport (L4) and finally application flows (L7), combined into conversations on each of the layers of the network protocol stack. Every conversation holds information about the source and destination endpoints (IP addresses, ports), timestamps and reassembled payloads of exchanged application messages.
- Data analysis — The analysis of each application conversation consists of identification of the application protocol, and extraction of application events, e.g., visited web pages, sent emails, DNS queries, etc., with proper application protocol dissector that yields sets of forensic artifacts.

A. Data Prepossessing

The First stage is executed on a set of independent *Reassembler* nodes. These reconstruct L7 conversations from the stream of captured packets which can originate from *PCAP files* or can be captured from *the live network interface*.

In the most common use-case, there are multiple source streams, i.e., a collection of PCAP files or direct network captures using PCAP-over-IP, which we want to analyze. Therefore to utilize all of the *Reassembler* instances, we have to split packets from each stream into a smaller sub-streams, which will be distributed among *Reassembler* instances. For this split, we cannot use a naive method such as *Round Robin*, because *Reassembler* nodes operate independently of each other and to fully reconstruct L7 conversation a particular *Reassembler* have to obtain all the pieces of that particular L7 conversation. In case we would use *Round Robin*, a situation could occur when a half the packets from one L7 conversation would end up in one *Reassembler* node and the second half in another; both nodes would have incomplete data and none of them would be able to reconstruct the conversation entirely.

Solution to this problem is another type of nodes called *L4 Load Balancer*, which will be positioned in front of the *Reassembler* nodes. *L4 Load Balancer* will extract source and destination IP addresses and ports and transport protocol from each packet of the source stream, and use them to decide to which instance of the *Reassemblers* should forward the packet based on its context. This way, all packets of a particular L7 conversation will always be forwarded to only one *Reassembler* instance. The reconstructed L7 conversation will be then stored in a distributed database, ready to be retrieved in the second stage of the execution.

B. Data Analysis

In the *second stage*, a subset of reconstructed L7 conversations is retrieved from the distributed database and delivered to the *Application protocol dissector* nodes. For every L7 conversation, *Application protocol dissector* nodes identify used application protocol and use a proper dissector module dedicated to the processing of a single application protocol such as HTTP, SMTP or DNS, to extract application protocol messages from this L7 conversation. Obtained data are stored back into the distributed database. Processing of application messages is under normal circumstances possible only with unencrypted network communication. From SSL/TLS communication which encapsulates application protocols such as HTTP, we can extract only readable portions of this data such as the server's SSL certificate. Possible ways to decrypt and subsequently, parse an SSL/TLS communication is to own a private key of a given SSL/TLS server or to deploy an SSL/TLS intercepting proxy [19].

Each instance of a node acts as an individual actor, communicating with other actors by message passing. Thanks to this design, we can distribute the computation across multiple machines maintaining the linear scalability.

V. PRELIMINARY EVALUATION

Our prototype implementation is based on C# actor system library *Akka.NET*. For testing and performance benchmarking, we have implemented two modes of operation:

1) *Offline* — isolated execution which combines the functionality of a single *L4 Load Balancer*, *Reassembler* and *Application protocol dissector* node inside one system's process. Therefore, no inter-actor message serialization occurs because data reside only inside a shared memory.
2) *Online* — distributed execution spanning across multiple cluster nodes. The inter-actor message serialization is required, because data leave shared memory space, and are serialized in-order to be delivered to a neighboring node.

Additionally, for proof-of-concept benchmarking, the functionality of *Application protocol dissector* nodes was included inside *Reassembler* nodes to eliminate distributed database as a middleman between them. In the following measurements, we focus on a raw network capture's processing performance of the so-far naive implementation. Currently, our prototype implementation supports the dissection of two application protocols (DNS and HTTP).

We have measured the preliminary performance, of the implementation, on different hardware configurations:

- *Workstation* — Intel i7-5930K 4.3 GHz, 12 cores, 64 GB RAM, 512 GB SSD
- *Mini-cluster* — 4x servers with Intel Xeon E5520, 2.26 GHz, 8 cores, 48 GB RAM, 1 TB SSD, 1 Gbps network

We used a public data set of M57-Patents Scenario [20], that consists of real-world data captured over a month.

We merged all network traces into one PCAP file of roughly 4.8 GB and 5,707,845 frames. One large PCAP file simulates our use-case of streamed-in communication that needs to be load-balanced from a single node.

	Workstation [Mbps]	**Mini-cluster** [Mbps]
PCAP file read	5103	5719
Packet parsing	3853	1679
L7 Conversation tracking	942	380
HTTP & DNS extraction	880	358

TABLE I. PROCESSING SPEEDS OF OUR OFFLINE TEST SCENARIO

Reassemblers count	**One** [Mbps]	**Two** [Mbps]	**Three** [Mbps]
HTTP & DNS extraction	233	407	453

TABLE II. PROCESSING SPEEDS OF OUR ONLINE TEST SCENARIO MEASURED ON MINI-CLUSTER

Workstation [Mbps]	**Mini-cluster** [Mbps]	**Netfox** [Mbps]	**Wireshark** [Mbps]	**NetworkMiner** [Mbps]
880	358	65.6	73.4	15.8

TABLE III. PROCESSING SPEEDS OF COMMONLY USED NETWORK FORENSIC TOOLS MEASURED ON WORKSTATION

We started with measurements in an *offline* mode on a single machine, firstly with a PCAP file parsing operation and incrementally added consequent operations and measured processing speeds, as Table I describes. Preliminary evaluation suggests that the *raw speed* of roughly 3.8 Gbps, for file reading and packet parsing is sufficient. The process of conversation tracking that segregates IP flows by packet source and destination IP addresses, ports and transport protocol type with additional heuristics [18], that also reassembles TCP/UDP streams, is computationally heavier, reaching "only" 942 Mbps, and is about 4x slower than only read and parsing. With added HTTP & DNS dissection, performance slightly decreased further down to 880 Mbps.

The *CPU frequency* plays a very important part, that can be observed if we compare our *Workstation* with *Mini-cluster* — 880 Mbps vs. 358 Mbps. All other components except CPUs are otherwise roughly comparable as we can see by comparing the speed of "PCAP file reading."

The scalability is described in Table II that show performance in *online* mode. The solution was deployed on *Mini-cluster*. The first node was reading the captured communication from a PCAP file and load-balancing it to the rest that reassembled it and dissected *HTTP* and *DNS* artifacts. In the measurements, we can see an increase in performance with each added *Reassembler*. Nevertheless, further optimization is required to achieve linear scalability.

We compare our solution running in the *offline* mode with commonly used network forensic tools in Table III. Our solution is an order of magnitude faster while delivering a comparable amount of results.

VI. CONCLUSION

In this research, we proposed a system for distributed real-time forensic network traffic analysis up to the application layer capable of large-scale communication processing. We intend to create a system based on the actor model that scales linearly and is hardware independent. The implementation environment of .NET Core framework and C# language enables rapid development compared to C/C++ that is used by VAST and Moloch. Also, our solution is multiplatform and easily staged with Docker Swarm. Therefore, the deployment of the entire computation cluster is reduced to one command. The solution is distributed under MIT License and hosted as an open-source project on GitHub — https://github.com/nesfit/NTPAC.

Shortly, we plan to measure the performance of our solution using data from real-world cases. Because of legal reasons, deployment to public cloud infrastructure is out of the question. Therefore, we need to build a private one that consists of nodes with high CPU frequencies and 10 Gbps network interfaces. Additionally, we need to profile and optimize processing mechanisms; expand the set of protocols supported by application protocol dissectors and add support for tunneling mechanisms.

VII. ACKNOWLEDGEMENT

This work was supported by BUT project "ICT Tools, Methods and Technologies for Smart Cities" (2017-2019), no. FIT-S-17-3964.

REFERENCES

[1] N. Beebe, "Digital forensic research: The good, the bad and the unaddressed," in IFIP International Conference on Digital Forensics. Springer, 2009, pp. 17–36.

[2] E. S. Pilli, R. C. Joshi, and R. Niyogi, "Network forensic frameworks: Survey and research challenges," digital investigation, vol. 7, no. 1-2, 2010, pp. 14–27.

[3] M. Vallentin, "Scalable network forensics," Ph.D. dissertation, UC Berkeley, 2016.

[4] S. Rekhis, J. Krichene, and N. Boudriga, "Digfornet: digital forensic in networking," in IFIP International Information Security Conference. Springer, 2008, pp. 637–651.

[5] A. Almulhem and I. Traore, "Experience with engineering a network forensics system," in International Conference on Information Networking. Springer, 2005, pp. 62–71.

[6] W. Wang and T. E. Daniels, "A graph based approach toward network forensics analysis," ACM Transactions on Information and System Security (TISSEC), vol. 12, no. 1, 2008, p. 4.

[7] N. L. Beebe and J. G. Clark, "A hierarchical, objectives-based framework for the digital investigations process," Digital Investigation, vol. 2, no. 2, 2005, pp. 147–167.

[8] S. Perumal, "Digital forensic model based on malaysian investigation process," International Journal of Computer Science and Network Security, vol. 9, no. 8, 2009, pp. 38–44.

[9] W. Halboob, R. Mahmod, M. Abulaish, H. Abbas, and K. Saleem, "Data warehousing based computer forensics investigation framework," in 2015 12th International Conference on Information Technology-New Generations (ITNG). IEEE, 2015, pp. 163–168.

[10] Bro, cited February 2019. [Online]. Available: https://www.zeek.org/

[11] Vast, cited February 2019. [Online]. Available: http://vast.io/

[12] Moloch, cited February 2019. [Online]. Available: https://molo.ch/

[13] W. Ren and H. Jin, "Distributed agent-based real time network intrusion forensics system architecture design," in Advanced Information Networking and Applications, 2005. AINA 2005. 19th International Conference on, vol. 1. IEEE, 2005, pp. 177–182.

[14] D. Wang, T. Li, S. Liu, J. Zhang, and C. Liu, "Dynamical network forensics based on immune agent," in Natural Computation, 2007. ICNC 2007. Third International Conference on, vol. 3. IEEE, 2007.

[15] S. Khan, A. Gani, A. W. A. Wahab, M. Shiraz, and I. Ahmad, "Network forensics: Review, taxonomy, and open challenges," Journal of Network and Computer Applications, vol. 66, 2016, pp. 214–235.

[16] M. I. Cohen, D. Bilby, and G. Caronni, “Distributed forensics and incident response in the enterprise,” in Digital Investigation, 2011.

[17] J. Pluskal, “Framework for captured network communication processing,” Ph.D. dissertation, Diploma thesis, FIT VUT v Brne, 2014, 2014.

[18] P. Matoušek, J. Pluskal, O. Ryšavỳ, V. Veselỳ, M. Kmet’, F. Karpíšek, and M. Vymlátil, “Advanced techniques for reconstruction of incomplete network data,” in International Conference on Digital Forensics and Cyber Crime. Springer, 2015, pp. 69–84.

[19] S. Davidoff and J. Ham, Network Forensics: Tracking Hackers through Cyberspace. Prentice Hall, 2012.

[20] M57-Patents Scenario, cited February 2019. [Online]. Available: https://digitalcorpora.org/corpora/scenarios/m57-patents-scenario