

TARZAN: An Integrated Platform for Security Analysis

Marek Rychlý, Ondřej Ryšavý
Brno University of Technology

Faculty of Information Technology, Department of Information Systems
IT4Innovations Centre of Excellence
Brno, Czech Republic
Email: {rychly, rysavy}@fit.vutbr.cz

Abstract—In this paper, we present the TARZAN platform, an integrated platform for analysis of digital data from security incidents. The platform serves primarily as a middleware between data sources and data processing applications, however, it also provides several supporting services and a runtime environment for the applications. The supporting services, such as a data storage, a resource and application registry, a synchronization service, and a distributed computing platform, are utilized by the TARZAN applications for various security-oriented analyses on the integrated data ranging from an IT security incident detection to inference analyses of data from social networks or crypto-currency transactions. To cope with a large amount of distributed data, both streamed in real-time and stored, and for the need of a large scale distributed computing, the platform has been designed as a big data processing system ensuring reliable, scalable, and cost-effective solution. The platform is demonstrated on the case of a security analysis of network traffic.

I. INTRODUCTION

THE ABUNDANCE of data sources and the exponential growth in the volume they produce represents new challenges for many traditional ICT fields. Digital forensics and security incident analysis is not an exception. Every day, security analysts and investigators face the problem of insufficient tool support. The roots of this problem lie in the fact that this dramatic change in the heterogeneity and volume of data makes the existing methods obsolete.

Traditional workflow of digital forensic consists of the well-defined procedure of data identification, acquisition, preservation, analysis, and reporting. This workflow was devised and refined in the 1990s when the environment regarding computing technology and software was rather uniform. Also, the amount of data that needs to be processed was from our perspective rather small. For most cases, it was possible to perform all above-mentioned steps using a single forensic workstation. Because of the rapid technology advances in the ICT, this is no longer true. Not only the increasing amount of data caused by the drop of storage cost and dissemination of broadband connectivity represents the challenge for digital

This work was supported by Ministry of Interior of the Czech Republic project “Integrated platform for analysis of digital data from security incidents” VI20172020062; Ministry of Education, Youth and Sports of the Czech Republic from the National Programme of Sustainability (NPU II) project “IT4Innovations excellence in science” LQ1602; and by BUT internal project “ICT tools, methods and technologies for smart cities” FIT-S-17-3964.

forensics. Often even the data acquisition phase is difficult to achieve with the existing tools and considering the usual methods of creating the forensic image of the disk drive. Completing this operation for nowadays common terabyte hard drive lasts several hours.

Investigators also need to face the problem of high diversity of computing devices. Smartphones, tablets and other connected smart devices massively penetrate the market. Cloud services are another emerging technology that changes the requirements on the digital forensics methods. All of this means that classical approach represented by well-defined workflow and considering the use of a single forensics workstation cannot meet the current demands. In many cases, the amount of data that to be processed exceeds several terabytes. Also, some forms of cyber crime comprise of the combination of several sophisticated techniques, and for their investigation, it is necessary to process and correlate information from several big datasets. To cope with this problem, several researchers suggested to apply big data approach, e.g. [1], and this field has become an active research area.

In this paper, an integrated platform for analysis of digital data from security incidents (a TARZAN platform) is proposed to address the issues mentioned above. The platform allows to gather, store, and process digital forensic data as big data to perform various security-oriented analyses that range from an IT security incident detection to inference analyses of data from social networks or crypto-currency transactions.

The paper is organized as follows. Section II discusses related work on data security analysis and processing platforms. In Section III, we provide a case study of a PCAP analysis tool utilizing the proposed platform for real-time security analysis of network traffic. Section IV introduces the TARZAN platform and describes its architecture and core services. The results of the case study implementation on the platform are discussed in Section V. Finally, we draw conclusions in Section VI.

II. RELATED WORK

Several approaches were already proposed to perform security, forensic, and inference analyses. Because conventional technologies are not always adequate to support long-term, large-scale analytics [2], big data approaches to the digital

forensics started to emerge addressing their own challenges (see [3], [1]). However, the most of the existing approaches focused on particular selected topics of IT security, related often only to networking security, rather than providing a general framework to support and integrate various forensic data to analyse them and their inferences in a broader context, such as in the cases of [4] and [5].

In [4], a digital forensic data reduction process were proposed based on a selective imaging, to speed up the forensic process by locating evidences, or by providing examiners with a quick understanding of the data to enable a better focus for full analysis (e.g., for a cross-device or cross-case analysis). Although the proposed process is general enough to support the examination of various types of the stored big data, it is not designed for custom autonomous big data analyses.

Feature Collection and Correlation Engine (FCCE, [5]) was introduced to find correlations across a diverse set of data types spanning over large time windows with very small latency and with minimal access to raw data. The engine entailed a complete framework for ingesting, aggregating, storing, as well as retrieving big data, by implementing feature extraction, aggregation, storage, and retrieval APIs, respectively. It was applied in IT security to detect fluxing domain names and identify persistent threat infections. However, the engine did not provide an implementation platform to build system utilizing the implemented APIs.

Network forensic analysis, which is the subject of the case study presented in this paper to demonstrate the TARZAN platform (see Section III), comprises of methods for capturing, collecting and analysing network data for information gathering, evidence identification, or security incident investigation. A new generation of Internet services opens space for new cybercrime activities. Security analyst and Law Enforcement Agency officers have to act accordingly to detect unlawful or unauthorized activities efficiently. The investigation is not possible without the tool support. While technology advances provide hardware technology able to capture communication at speeds that match current wire speed the software equipment for analysis of captured traffic has difficulties with packet traces of several gigabytes.

Network forensic analysis methods were implemented in various tools. General purpose tools include network analysers (Wireshark, TCP dump), IDS systems (Snort, Bro), fingerprinting tools (Nmap, p0f), and tools to identify and analyse security threats.

Tools dedicated to network forensic analysis implement specific features to aid investigation process. They can capture an entire network traffic and allow an investigator to analyse it and reconstruct the communication. Several widely used open source tools exist. In the following, we briefly overview three freely available tools that employ the typical features. Survey of network forensic tools can be found in [6].

PyFlag is a general purpose forensic package which can be used as disk forensics, memory forensics, and network forensics tool. This tool was developed by M. Cohen of Australian Federal Police in 2005 [7]. PyFlag is designed around the

Virtual File System concept. For each supported data source a specific loader is implemented. To deal with PCAP files, the PCAP filesystem loader opens PCAP file, parses and dissects individual packets up to lower layer protocols, collects related TCP packets into streams and finally applies higher level protocol dissectors. A forensic investigator is usually interested in high-level information that can be extracted from the communication. PyFlag enables to reassemble the content of communication, e.g., web pages, email conversation, etc.

Network Miner¹ is an open source tool that integrates packet sniffer and higher-layer protocol analysers into a tool for passive network traffic monitoring and analysis. Because captured traffic can be processed in the same way, Network Miner is also a valuable tool for network forensics analysis. Network Miner offers several useful features, such as the possibility of operating system identification, traffic classification, and reassembling the transferred files for HTTP, FTP, TFTP and SMB protocols.

Xplico² is a modular tool aimed at the reconstruction of the data content carried in the network traffic. The software consists of the input module handling the loading source data, decoding module equipped with protocol dissectors for decoding the traffic and exporting the content, and the output module organizing decoded data and presenting them to the user. Contrary to PyFlag and NetworkMiner, Xplico is not a typical desktop application but it is deployed as a server service with the web-based interface. The authors claim the possibility to analyze large PCAP files of many gigabytes. Because the Xplico design is a classical client-server architecture, the performance of the tool is limited by the hardware configuration of the server running the Xplico backend.

To analyse the network traffic as big data, a scalable internet traffic analysis system was presented in [8]. The system, which was able to process multi-terabytes libpcap dump files, utilized Apache Spark for data processing to analyse captured transmitted data and protocol fields. Unfortunately, the system did not allow to integrate non-network data and perform the analyses of the network data in broader contexts.

Another approaches to the network traffic security big data analysis were presented in [9], [10], [11], and [12]. Apache Metron [13] and Apache Spot [14] projects are more interesting. They try to form general frameworks for security analyses of IT threats, secondary processing also firewall and application logs, emails, intrusion-detection reports, etc. However, analogously to the first case, also the all of these approaches were focused primarily and narrowly on the network data and unable to find correlations with other forensic data or to provide a comprehensive platform for big data forensics.

III. CASE STUDY: PCAP ANALYSIS

Digital investigators process network traffic as a source of evidence in many types of computer crimes. Captured traffic can be analyzed to obtain the content and also to show

¹<http://www.netresec.com/?page=NetworkMiner>

²<https://github.com/xplico/xplico>

the actions taken by the offender. Network traffic can also be important for corroborating evidence. Obtaining network traffic as a source of evidence is usually more complicated than other digital evidence. Transmitted data are only available on the network device for a limited amount of time. Inappropriate collection method can lead to data corruption or incomplete capture. As messages exchanged between applications are segmented into many pieces, it is important to gather all relevant packets and be able to combine them again into data streams. When collecting data on shared links, there may be a huge amount of traffic from which only a fraction is relevant to the investigation. Moreover, many different protocols are in use which requires applying corresponding decoding algorithms. Although existing tools for information security can be adapted for a forensics investigation, they usually lack features for evidence collecting and preservation. For the forensic investigation, there are two important activities, namely examination and analysis [15]. The examination is characterized by the mostly automatic data processing that ends with a collection of relevant data extracted from the data source. The analysis follows examination, and it is often a manual or more interactive activity that interprets the significance and meaning of the extracted data. Also, data correlation, finding links and patterns in the extracted data is the desired result of the analysis.

From the examination viewpoint, the important features of network forensic tools are as follows:

- *Flow reconstruction.* Because network conversation is split into many packets exchanged by communicating applications, the first step of data examination is to combine these fragments to form flows again.
- *Protocol identification.* Network communication is governed by protocols. There are many protocols in the Internet communication. The ability to identify which protocol was used to data exchange is crucial for further processing. Protocol identification is difficult for encrypted traffic where traditional pattern based methods may be less accurate results.
- *Protocol decoding.* To understand the communication we often need to extract data from protocol header fields and data payload. Network forensic tools support a wide variety of protocols. Sometimes these decoders can identify anomalous packets that do not conform to the protocol specification.
- *Data reduction.* Not all data needs to be analyzed. Data reduction can filter out uninterested data. The filter applied depends on the information obtained from the protocol decoding step. We can be, for instance, interested only in Web traffic.
- *Data recovery.* If communication is not encrypted, the communication payload is visible. This gives us the possibility to recover digital objects from the communication such as web pages, images, e-mail messages.
- *Pattern search.* The common investigative approach is to search for occurrences of known patterns, e.g., email

addresses, keywords, etc. Pattern search in network traffic needs to consider specifics of various protocols, such as encoding, compressing, etc.

Forensic data analysis can involve different methods and procedures. The following techniques are commonly applied:

- Developing the *timeline* from significant events offers investigators a high-level view on the extracted data. Different kinds of communication can contribute to timeline by various events, such as e-mail delivery, web search, file download, etc.
- The *temporal analysis* aims to identify patterns or anomalies that are often processed by the further and deeper analysis. For instance, we are seeking for the periods of peak data transfer or occurrences of an unusual protocol.
- The *relation analysis* provides links among different entities. Relations can be analyzed on different layers, linking devices, services, or end-users
- *Classification* methods assign extracted data to different classes according to the predefined criteria, such as system traffic, web traffic, suspicious traffic, malware related traffic, etc.
- *Clustering* techniques can deal with a lot of entities by grouping them according to some essential characteristics. Often these methods do not require learning and thus are easily applicable.
- *Correlation* is a statistical technique that can identify the relation between different entities. It is, for instance, possible to identify the same activity recorded in various data sources.

Digital investigation is a time-consuming and labor-intensive process. Thus, there is a strong emphasis on using tools that can reduce the examination time and improve the automatization of analysis activities. In the next section, we will show, how the proposed platform can achieve both requirements. First, examination time can be reduced by allocating more computation nodes. Second, some analysis can be automated by applying machine learning algorithms.

The complex PCAP analysis requires processing of a huge amount of data. The processing must be done both in real-time to detect security incidents or to perform security audits, and later on large stored datasets to answer queries of an analyst. As the such processing is difficult to do by conventional centralized approaches in highly scalable, high-throughput, and fault-tolerant way [2], the PCAP analysis tool will be implemented on the TARZAN platform.

IV. THE TARZAN PLATFORM

To ensure communication of TARZAN applications and provide them with required services and a runtime environment, the TARZAN platform consists of three core components, namely, *Platform Bus* which implements a distributed communication bus for the applications and the components, *Platform Storage* which provides a distributed storage service (NoSQL databases, distributed filesystem, resource registries, etc.), and *Platform Computation* component to provide the

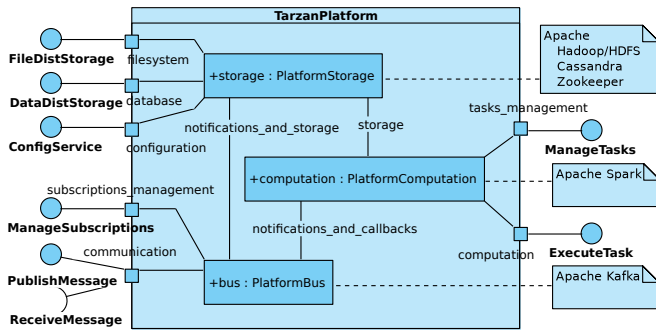


Fig. 1. The TARZAN platform architecture.

runtime environment for distributed computation tasks of TARZAN applications.

In Figure 1, architecture of the TARZAN platform is modelled in an UML composite structure diagram. Each of the three core components provides its services to TARZAN applications by the platform's external interfaces. Moreover, the components communicate and cooperate inside the platform. The individual components are described in the following sections.

A. Platform Bus

The main goal of the platform bus core component is to enable asynchronous communication of other TARZAN components. More specifically, the platform bus implements a *publisher-subscriber* communication model based on *message queues*. A client is able to publish messages to particular topics acting as a producer, or to subscribe to receive messages of particular topics as a consumer (see the corresponding interfaces in Figure 1). The platform bus guarantees delivery of the published messages to their appropriate consumers.

The communication via the bus is utilized by both external TARZAN applications and the core TARZAN components. In the first case, the applications can connect themselves to various data sources to ingest sensor data, events, logs, etc.; interconnect their components into data processing topologies to perform data parsing, normalizing, validating, marking, enrichment, etc.; and consume or feed data from/into the platform storage components. In the second case, the platform bus helps the other core components to send/receive their data, for example, to store the transmitted data into the platform storage and deliver the storage update notifications, or to deliver input data and pass output data of tasks of the platform computation including callbacks.

To achieve high-throughput message passing in highly scalable distributed environments, the platform bus is based on *Apache Kafka* [16]. In Kafka, messages are communicated in topics. Each topic, as a general category of particular messages, consists of multiple partitions (queues). A producer submits a message to a particular topic (or topics) where in each topic, the message is assigned to a particular single partition (automatically for load-balancing or as required by the producer). A consumer can belong to a particular consumer

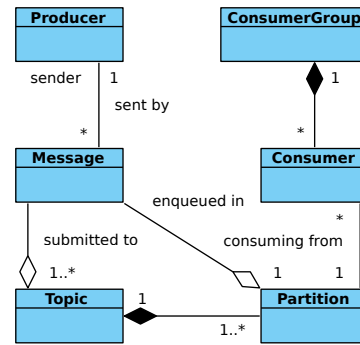


Fig. 2. A conceptual model of basic entities in Apache Kafka.

group and subscribes to one or more topics. In each of the subscribed topics, the consumer has assigned particular partition for exclusive reception. For relationships of these concepts, see Figure 2.

In TARZAN, Kafka's concepts of a message, topic, partition, producer, consumer, and consumer group are utilized for consuming data sources and communication with computation tasks as follows.

1) *Broadcasting from data sources*: A data source (producer) submits data (a message) of a particular type (topic) under the data source's identification (partition). A subscriber (consumer) listens to a particular topic and a particular partition, that is for messages of the particular type from the particular data source. A message will be received by (broadcasted to) all subscribed consumers in different consumer groups.

- Messages = data produced by the sources.
- Topics = individual data source types (e.g., PCAP).
- Partitions = particular data sources (e.g., a sensor monitoring a network traffic on a particular network interface).
- Consumer groups = subscribers for data produced by a particular data source (e.g., a component for processing/analysing/storing PCAPs).

2) *Load-balancing of data processing tasks*: A client (producer) submits a task invocation (message) to a particular service (topic) without any partition (it will be assigned automatically by Kafka for load-balancing). In the case of a request-reply task invocation, the message should contain also the client's identifier which will be utilized for the callback (a particular partition name in "callback" topic).

- Messages = task invocations including data payloads and callback addresses if needed.
- Topics = names of individual services (e.g., PCAP Analyzer).
- Partitions = individual instances of a particular service (e.g., a particular process of the PCAP Analyzer).
- Consumer groups = single-member groups representing the instances as above.

3) *Delivery of the tasks' replies*: A particular service task instance (producer) submits a reply/result (message) to the previously received task invocation as a callback. The reply (message) will be delivered to a particular client who sent the

task invocation (to his partition in “callback” topic).

- Messages = replies/results to the previously submitted task invocations.
- Topics = a single topic with name "callback" only.
- Partitions = one partition for each individual client.
- Consumer groups = single-member groups representing the clients as above.

B. Platform Storage

While the platform bus described in the previous section is necessary for data processing, the platform storage implements the data persistence in distributed environments. The distributed data storage is the necessary requirement of distributed data processing to be able to scatter and deliver data across processing nodes. Three types of data storage services are supported: a distributed filesystem, a distributed database, and a distributed and synchronized configuration service (see the corresponding interfaces in Figure 1).

The platform storage services are utilized by both external TARZAN applications to provide a shared storage and by the core TARZAN components to store the platform runtime data. In the second case, the storage services are utilized for a resource registry of various resources accessed and manipulated by the platform (e.g., topic and partition names for the platform bus, or declarations and definitions of tasks in the platform computation components).

For the distributed filesystem and the distributed and synchronized configuration service, Hadoop Distributed File System (HDFS) from *Apache Hadoop* [17] and *Apache Zookeeper* [18] were adopted, respectively. Both software products are widely utilized in the TARZAN platform and well-integrated with other components. For example, the platform bus based on Apache Kafka is utilizing Zookeeper for message queue management and the platform computation component based on Apache Spark is utilizing HDFS for a data storage and Hadoop for a cluster management.

Although the distributed database service is not designated for a particular NoSQL database, *Apache Cassandra* [19] is the preferred database in the TARZAN platform. The main reason for this preference is a perfect integration of Cassandra with the rest of the software stack (e.g., well-established Cassandra-Spark and Cassandra-Kafka connectors). Moreover, Cassandra provides an optimal storage for large sensor data [20].

C. Platform Computation

To support distributed computing on data communicated and stored in the TARZAN platform, the platform computation core component is provided. The component allows TARZAN applications to run tasks, e.g., to process (normalize/aggregate), enrich, label, combine, etc. the data and to utilize other TARZAN components.

Tasks for the platform computation component are registered by external application components and then they can be executed by TARZAN applications (for the corresponding interfaces, see Figure 1) as demonstrated in Figure 3 to perform malware or data-breach detections, or to analyse

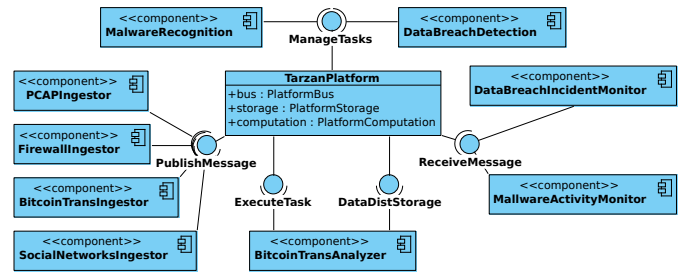


Fig. 3. An example of external application components utilizing the TARZAN platform (the ingestors on the left side are feeding data to the platform, computation tasks and an application on the top and bottom are processing the data, and the monitors on right side are passing results to clients).

Bitcoin transaction based on capture network traffic, firewall logs, Bitcoin blockchain, and social network profiles.

As the most of the use-cases for data processing in the TARZAN platform operate on big data (in the sense of data characterized by four Vs: volume, variety, velocity, and value [21]), the platform computation tasks must be able to do big data processing. The applications need to process both data streams and data batches (e.g., to do a real-time analysis of firewall logs and to execute on-demand tasks, respectively). Therefore, *Apache Spark* [22] has been selected as the implementation technology for the platform computation component and its tasks, as it supports both the stream and batch processing of big data.

For the batch data processing in Spark, computation tasks can utilize a data abstraction called *Resilient Distributed Dataset* (RDD) to execute various parallel operations on a Spark cluster and to gather resulting data in shared broadcast variables and accumulators provided by Spark on the cluster’s nodes. In the case of the stream data processing, Spark Streaming provides computation tasks with *Discretized Stream* (DStream) abstraction where each stream is represented by a continuous series of RDDs that is divided into micro-batches and processed by the tasks in the similar way as in the batch processing above. Because DStreams follow the same processing model as batch systems, the two can naturally be combined [23] and the platform computation component and its tasks can implement identical algorithms for both the stream and batch processing.

V. EVALUATION

The TARZAN platform has been evaluated by means of the PCAP analysis case study described in Section III. A corresponding TARZAN application has been implemented to read and analyse data of network traffic monitoring stored in the PCAP format. After the PCAP data are read from input PCAP files or real-time generated by network traffic monitoring tools, they are transferred (including their related meta-data) via platform bus for a primary analysis by tasks of platform computation. The tasks also ensure that both the input data and the output primary analysis results are stored in platform storage. Afterwards, a secondary analysis can

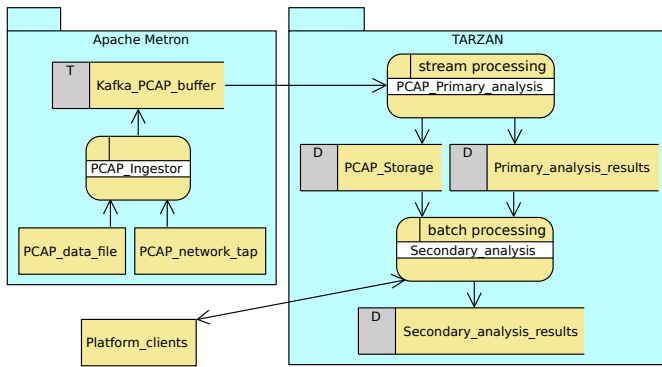


Fig. 4. Architecture of the PCAP Analysis tool with data-flows (including processes, data storages, and external data sources and entities).

be executed on the stored data and the previous results to perform various security and forensic analyses, e.g., to detect communication patterns, apply clustering methods for data mining, etc.

The primary analysis is operating on continuously incoming data and the primary analysis tasks implement real-time stream processing to quickly extract traffic basic features such as source and destination IP addresses and port numbers, defragment fragmented packets into network flows, analyse flow properties, application protocols, etc. These tasks utilize the Spark Streaming extension of the core Apache Spark API to process DStreams. In Spark, tasks are scalable, high-throughput, fault-tolerant, so the ability to process the incoming live data in real-time can be improved, if necessary, by an appropriate cluster configuration and the application deployment. However, the primary analysis must perform only basic analytical tasks.

Contrary to the primary analysis which employs real-time stream processing, the secondary analysis can implement a batch processing of the previously stored data and the primary analysis results. Therefore, the stored inputs can be represented as RDDs and processed by means of Spark RDD API, Spark SQL, and also machine learning algorithms provided by Spark's machine learning library (MLlib) can be applied. The secondary analysis is executed on demand as required by the platform's client applications, e.g., to provide data for visualisations, analyse network communications related to security incidents under investigation, or related to cryptocurrency transactions or malware activities.

The overall architecture of the PCAP Analysis tool is depicted in Figure 4. To feed input PCAP data into the system, several modules were adopted and adapted from the *Apache Metron* project [13], namely: *metron-sensors*, *metron-pcap*, and *metron-api*. In the first module, Apache Metron brings into TARZAN the integration of *Data Plane Development Kit*³ (DPDK) probes for high speed packet capture and *Yet Another Flowmeter*⁴ (YAF) to processes packet data from

PCAP dumpfiles (as generated by *tcpdump* or *libpcap*). The next two Metron modules provide a topology for streaming network packets into HDFS and low-level analytics/filtering on the PCAP files in HDFS before they are submitted into a Kafka message queue acting as a buffer for further processing. Then, a continuous stream processing in the primary analysis and an on-demand batch processing in the secondary analysis is performed by utilizing the TARZAN platform components as described above.

In comparison with the *Apache Metron* [13] or *Apache Spot* [14] discussed in Section II, the current implementation of the PCAP analysis tool in TARZAN provides the same basic functionality, however, it enables a better integration with the other TARZAN applications into a seamless security analysis framework where results of the PCAP analyses may contribute to various security investigations, e.g., to trace cryptocurrency transactions or malware activity.

In comparison with the existing approaches and the related work (see Section II), the TARZAN platform is a step further in the design and development of open forensic platform capable of processing big data. As we demonstrated in the PCAP analysis case study, our approach is compatible and easily integrated with other approaches to big data forensic. TARZAN applications can utilize HDFS as suggested in a conceptual model of big data forensics by Zawoad and Hasan [24]. Also a framework for the forensic analysis of big heterogeneous data presented by Mohammed et al. [25] can be realized using the TARZAN platform. Their framework has three layers that follow acquisition, examination, and analysis approach to extract metadata from acquired data sources and build a semantic web-based model for further analysis. While they do not specify the particular implementation of such system, the presented concepts are in accordance with the architecture of the TARZAN platform. Analogously, Irons and Lallie [26] discussed the shortcomings of the current analysis methods and suggested to use more intelligent techniques and demonstrated the possible application of artificial intelligence (AI) to computer forensics. The TARZAN platform can easily integrate the AI investigative methods because the underlying components provide rich libraries of various AI algorithms.

VI. CONCLUSION

In this paper, we have introduced a TRAZAN platform, an integrated platform for analysis of digital data from security incidents. The architectural design has been presented to explain which core component are available in the platform and which services are provided to TARZAN applications. The platform allows to gather, store, and process digital forensic data as big data to perform various security-oriented analyses.

As a sample case study, a PCAP analysis tool has been implemented on the platform utilizing the platform bus component to integrate individual modules, the platform storage component to store input data and analyses results, and the platform computation component to perform both stream and batch processing of big data.

³<http://dpdk.org/>

⁴<https://tools.netsa.cert.org/yaf/>

It has been concluded that the TARZAN platform constitutes an open forensic platform capable of processing big data and provides a sufficient framework for further integration of various existing approaches. The integration of various existing approaches and existing tools for forensic analyses as external TARZAN components and applications is a part of ongoing work.

REFERENCES

- [1] A. Guarino, *Digital Forensics as a Big Data Challenge*. Wiesbaden: Springer Fachmedien Wiesbaden, 2013, pp. 197–203. ISBN 978-3-658-03371-2. [Online]. Available: http://dx.doi.org/10.1007/978-3-658-03371-2_17
- [2] A. A. Cardenas, P. K. Manadhata, and S. P. Rajan, “Big data analytics for security,” *IEEE Security Privacy*, vol. 11, no. 6, pp. 74–76, Nov. 2013. doi: 10.1109/MSP.2013.138. [Online]. Available: <http://dx.doi.org/10.1109/MSP.2013.138>
- [3] H. V. Jagadish, J. Gehrke, A. Labrinidis, Y. Papakonstantinou, J. M. Patel, R. Ramakrishnan, and C. Shahabi, “Big data and its technical challenges,” *Commun. ACM*, vol. 57, no. 7, pp. 86–94, Jul. 2014. doi: 10.1145/2611567. [Online]. Available: <http://doi.acm.org/10.1145/2611567>
- [4] D. Quick and K.-K. R. Choo, “Big forensic data reduction: digital forensic images and electronic evidence,” *Cluster Computing*, vol. 19, no. 2, pp. 723–740, 2016. doi: 10.1007/s10586-016-0553-1. [Online]. Available: <http://dx.doi.org/10.1007/s10586-016-0553-1>
- [5] D. L. Schales, X. Hu, J. Jang, R. Sailer, M. P. Stoecklin, and T. Wang, “FCCE: Highly scalable distributed feature collection and correlation engine for low latency big data analytics,” in *2015 IEEE 31st International Conference on Data Engineering*, Apr. 2015. doi: 10.1109/ICDE.2015.7113379. ISSN 1063-6382 pp. 1316–1327. [Online]. Available: <http://dx.doi.org/10.1109/ICDE.2015.7113379>
- [6] E. S. Pilli, R. Joshi, and R. Niyogi, “Network forensic frameworks: Survey and research challenges,” *Digital Investigation*, vol. 7, no. 1–2, pp. 14–27, 2010. doi: <https://doi.org/10.1016/j.diin.2010.02.003>. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1742287610000113>
- [7] M. I. Cohen, “Pyflag: An advanced network forensic framework,” in *Proceedings of the 2008 Digital Forensics Research Workshop*. DFRWS, Aug. 2008. [Online]. Available: <http://www.pyflag.org>
- [8] A. Lukashin, L. Laboshin, V. Zaborovsky, and V. Mulukha, *Distributed Packet Trace Processing Method for Information Security Analysis*. Cham: Springer International Publishing, 2014, pp. 535–543. ISBN 978-3-319-10353-2. [Online]. Available: http://dx.doi.org/10.1007/978-3-319-10353-2_49
- [9] M. Wullink, G. C. M. Moura, M. Muller, and C. Hesselman, “ENTRADA: A high-performance network traffic data streaming warehouse,” in *NOMS 2016 - 2016 IEEE/IFIP Network Operations and Management Symposium*, Apr. 2016. doi: 10.1109/NOMS.2016.7502925 pp. 913–918. [Online]. Available: <http://dx.doi.org/10.1109/NOMS.2016.7502925>
- [10] M. Aupetit, Y. Zhauniarovich, G. Vasiliadis, M. Dacier, and Y. Boshmaf, “Visualization of actionable knowledge to mitigate DRDoS attacks,” in *2016 IEEE Symposium on Visualization for Cyber Security (VizSec)*, Oct. 2016. doi: 10.1109/VIZSEC.2016.7739577 pp. 1–8. [Online]. Available: <http://dx.doi.org/10.1109/VIZSEC.2016.7739577>
- [11] N. Promrit and A. Mingkhwan, “Traffic flow classification and visualization for network forensic analysis,” in *2015 IEEE 29th International Conference on Advanced Information Networking and Applications*, Mar. 2015. doi: 10.1109/AINA.2015.207. ISSN 1550-445X pp. 358–364. [Online]. Available: <http://dx.doi.org/10.1109/AINA.2015.207>
- [12] L. He, B. Tang, M. Zhu, B. Lu, and W. Huang, *NetflowVis: A Temporal Visualization System for Netflow Logs Analysis*. Cham: Springer International Publishing, 2016, pp. 202–209. ISBN 978-3-319-46771-9. [Online]. Available: http://dx.doi.org/10.1007/978-3-319-46771-9_27
- [13] (2016) Apache Metron: Real-time big data security. [Online]. Available: <https://metron.incubator.apache.org/>
- [14] (2016) Apache Spot (incubating): A community approach to fighting cyber threats. [Online]. Available: <https://spot.incubator.apache.org/>
- [15] Eoghan and Casey, “Network traffic as a source of evidence: tool strengths, weaknesses, and future needs,” *Digital Investigation*, vol. 1, no. 1, pp. 28–43, 2004. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1742287603000033>
- [16] (2016) Apache Kafka: A high-throughput distributed messaging system. [Online]. Available: <https://kafka.apache.org/>
- [17] (2014) Welcome to Apache Hadoop! [Online]. Available: <https://hadoop.apache.org/>
- [18] (2010) Apache ZooKeeper. [Online]. Available: <https://zookeeper.apache.org/>
- [19] (2016) Apache Cassandra. [Online]. Available: <https://cassandra.apache.org/>
- [20] J. S. van der Veen, B. van der Waaij, and R. J. Meijer, “Sensor data storage performance: SQL or NoSQL, physical or virtual,” in *2012 IEEE Fifth International Conference on Cloud Computing*, Jun. 2012. doi: 10.1109/CLOUD.2012.18. ISSN 2159-6182 pp. 431–438. [Online]. Available: <http://dx.doi.org/10.1109/CLOUD.2012.18>
- [21] J. Gantz and D. Reinsel, “Extracting value from chaos,” *ITC view*, vol. 1142, no. 2011, pp. 1–12, 2011.
- [22] (2016) Apache Spark: Lightning-fast cluster computing. [Online]. Available: <https://spark.apache.org/>
- [23] M. Zaharia, T. Das, H. Li, S. Shenker, and I. Stoica, “Discretized streams: An efficient and fault-tolerant model for stream processing on large clusters,” in *Proceedings of the 4th USENIX Conference on Hot Topics in Cloud Computing*, ser. HotCloud’12. Berkeley, CA, USA: USENIX Association, 2012.
- [24] S. Zawoad and R. Hasan, “Digital forensics in the age of big data: Challenges, approaches, and opportunities,” *2015 IEEE 17th International Conference on High Performance Computing and Communications (HPCC), 2015 IEEE 7th International Symposium on Cyberspace Safety and Security (CSS), and 2015 IEEE 12th International Conf on Embedded Software and Systems (ICES)*, pp. 1320–1325, 2015. doi: 10.1109/HPCC-CSS-ICES.2015.305
- [25] H. J. Mohammed, N. L. Clarke, and F. Li, “An automated approach for digital forensic analysis of heterogeneous big data,” *JDFSL*, vol. 11, no. 2, pp. 137–152, 2016. [Online]. Available: <http://ojs.jdfsl.org/index.php/jdfsl/article/view/410>
- [26] A. Irons and H. Lallie, “Digital forensics to intelligent forensics,” *Future Internet*, vol. 6, no. 3, pp. 584–596, 2014. doi: 10.3390/fi6030584. [Online]. Available: <http://doi.acm.org/10.3390/fi6030584>