

EventFlow: Network Flow Aggregation Based on User Actions

Petr Velan
CESNET z. s. p. o.
Prague, Czech Republic
petr.velan@cesnet.cz

Abstract—Network flow monitoring is being supplemented by an application flow visibility to provide more detailed information about a network traffic. However, the current concept of flows does not provide a mechanism to keep track of a semantic relations between individual flows that are created as a part of a single user action. We propose an extension to the flow measurement, called EventFlow, which allows to preserve relations between HTTP and DNS application flows that are a part of single user action, most typically browsing to a web page. We describe an architecture of the EventFlow extension and its limitations. A prototype implementation of the EventFlow is introduced and evaluated on a packet trace from an ISP network. We show that a significant number of flow records can be recognised as a part of a single user action.

I. INTRODUCTION

The growth of cloud-based services increases the importance of network monitoring. Information about network traffic behaviour can not only provide valuable information for performance optimisation of applications and infrastructure but also help to detect and mitigate attacks on applications and their users. To better facilitate these demands the network traffic monitoring solutions are starting to provide application visibility [1], [2].

Application flow monitoring parses data from application headers and adds application specific elements to flow records. This way the information from application level can be easily transferred to flow collectors, stored and utilised together with the information about the network communication. Current approach is to treat separate application protocols individually, e.g. develop an application processing module for each monitored protocol [3]. However, connections between different protocols are lost in this scenario. For example, when a user wants to access a web page, several different flows records are created. The DNS server must be contacted to resolve the hostname of the web page to an IP address. After the basic document is loaded, the user's browser automatically loads linked content such as images, cascading style sheets, and java script libraries. The generated requests are recorded as flows, however little relation between the flows is preserved.

Information about relations between individual flows can be useful in several scenarios. First, when an advertisement on some page contains malware, the original page can be determined using the relation and notified of the malicious content. Second, aggregates of the related flows can be created to simplify behavioural analysis of the network traffic. Moreover, the analysis can use the additional information to improve its

accuracy. Last, traffic classification engines can also benefit from having access to information about flow relations [4].

In this paper we present a flow monitoring extension, called EventFlow, which allows to keep track of relations between HTTP and DNS application flows. Information about flow relation is inserted to flow records to keep track of individual user actions i.e. events. We develop a prototype of the EventFlow extension and evaluate its properties on a network traffic trace from an ISP network. Results show that at least 10 % of HTTP and DNS flow records form more complex events. We believe that this is only a lower bound and that further improvements can be made to relate even more flows into events.

The rest of the paper is structured as follows. Related work is surveyed in Section II. We propose the architecture of EventFlow measurement in Section III. Section IV describes the implementation of the EventFlow prototype. Experimental evaluation of the EventFlow prototype is performed in Section V. The paper is concluded in Section VI.

II. RELATED WORK

Madhyastha and Krishnamurthy [5] propose a generic language for application-specific flow sampling. Their language allows applications to select flows with special properties so that the negative impact of sampling on these applications is minimised. This can be useful for intrusion detection systems or traffic classification applications. Although the goal of this work is different from ours, it also aims to improve the collected data so that traffic analysis applications can achieve higher accuracy.

Authors of [6] also focus on improving quality of sampled flow data. They show that the traffic classification accuracy can be increased using related sampling, which gives higher probability to connections that are a part of the same application. Authors propose to use a source IP address as a measure of relation between connection sessions.

Hu et al. [7] propose an entropy based aggregation system to mitigate an impact of DoS attacks and worm spreads on a network monitoring system. The main contribution of their approach is a flow key attribute selection algorithm that chooses key attributes by which the flows are aggregated. Two dimensional hash table is used to implement their approach. The aggregated flows are called metaflows. The main difference from EventFlow is that we label existing flows belonging to same user action, while the metaflow is a substitute flow for many flows created during an malicious network activity.




Fig. 1. Relations between HTTP and DNS requests and responses.

Dolberg et al. [8] introduce a multidimensional flow aggregation aimed to reduce the volume of collected data. Authors use tree structures for storing the data by chosen dimension such as IP addresses or ports. EventFlow proposed in our work might be used in this scenario to aggregate flows by the same events.

The usual approach to reduce the volume of collected data is to use sampling. Estan et al. [9] propose to use adaptive sampling rate to achieve highest possible accuracy within given data collection constraints. Their main contribution is a system for renormalisation of flow entries after the sampling rate was changed. Authors propose an extension to standard flow counting that increases the accuracy of the counters for sampled flow.

III. EVENTFLOW ARCHITECTURE

This section describes the architecture of the EventFlow monitoring. The goal is to label all flows that are the results of a single user action with the same event identifier (EID). For example, accessing <http://www.w3.org/> creates 1 DNS request, 37 HTTP requests, and 8 HTTPS requests. We aim to assign a single unique EID to the flows generated for all DNS and HTTP requests.

Four basic types of flows are recognised by the EventFlow, HTTP requests, HTTP responses, DNS requests and DNS responses. There are relations between these types of flows in the network traffic, as shown in Figure 1. When HTTP request to a new site is performed, the IP address of the site must be resolved first. Therefore, a DNS request is created. After the request is observed, a reply follows, which results in the relation 1). After the DNS reply arrives, the client knows the IP address of the server and makes the HTTP request, which creates the relation 2). An HTTP response follows the request, as indicated by the relation 3). The HTTP response can contain an HTML page which links to several additional resources such as external style sheets or images. The loading of these resources triggers more HTTP requests, resulting in relation 4). When these requests point to previously unresolved domains, new DNS requests are created, which introduces the relation 5).

We base the EventFlow architecture on the relations between the requests and responses. When an HTTP or a DNS flow is encountered, we must make sure that it is assigned the same EID as the related flows. Therefore, we create four sets of records: expected HTTP requests, expected HTTP responses, expected DNS requests and expected DNS responses. When processing an HTTP or a DNS flow, we add new record to the set or sets it relates to. Then, when a next flow is processed,

it is matched against appropriate expected set to see whether it is a part of existing event. If it is, an EID of the event is assigned to the flow record. For example, when a DNS response is encountered, a new record is put into the expected HTTP requests set (because of relation 2), see Figure 1). Then, when an HTTP request is processed, we check the expected HTTP requests set to see whether we are expecting this request based on a previous DNS response. If the request is matched, it is assigned the same EID as the DNS response.

TABLE I. MATCHED FLOW PROPERTIES.

	Source IP	Destination IP	Destination Port	URL	Domain	DNS Transaction ID
Expected HTTP Request	✓			✓	✓	
Expected HTTP Reply	✓	✓	✓			
Expected DNS Request	✓				✓	
Expected DNS Reply	✓	✓	✓			✓

Each of the sets of expected records uses different flow properties to match a flow record. When matching an HTTP request flow against the expected HTTP requests set, the source IP address of the flow must match as well as the requested domain on the URL, if available. Checking the source IP address ensures that flows from different hosts are not combined into a single event. A domain name is checked for the records that were inserted in the set when DNS reply was encountered. In case an HTTP response caused the record to be inserted, a full URL is available, not only a domain name. Replies are checked based on IP addresses and destination port. Source port is not checked since the services are expected to run on standard well-known ports. The DNS reply is also checked for transaction ID, which is a unique identifier tying the request and response. List of the used properties is provided in the Table I.

An expiration of the records from the expected sets must be ensured. When a record from any of the expected sets is matched, it is removed. However, many records are inserted that will not be ever matched. For example, when a DNS request is made to accommodate a different service than HTTP, the expected HTTP request might never appear. We need to free such records from the sets eventually. A timeout is used to keep the expected sets from being congested by redundant records. A timestamp is assigned to each record upon insertion to a set. Then, each time a set is searched, records older than the timeout are removed. The timeout should be as short as possible to avoid blending of several events. However, it should be at least as long as it takes to process the longest user action, which might be up to a couple of seconds in case of complicated queries to slow sites.

There are several caveats to our approach and some limitations of the architecture that should be addressed in the future. Our approach does not handle HTTP redirection codes, therefore the first request and HTTP 3xx redirection response are assigned different EID than the subsequent request to the resource. This problem can be rectified simply by adding a handler for the HTTP 3xx redirection responses that will put




Fig. 2. EventFlow prototype schema.

a new record with the redirect URL to the expected HTTP requests set.

Another limitation of our approach is that the URLs are only extracted from HTML documents. However, modern web sites often use a JavaScript code to request additional resources through the Ajax technique. Such requests cannot be easily matched to an event since it would require to reconstruct the complete web page and process the included JavaScript code, which is unfeasible for the flow monitoring system.

There are also several caveats that cannot be avoided. Some of the requested documents might be cached by the clients which would cause EventFlow to loose track of related URLs. However, cached DNS queries are of no consequence to the EventFlow since no traffic is generated for them and no information about a flow relation is lost. And last, the growing deployment of HTTPS reduces the usefulness of the EventFlow for HTTP protocol. Nevertheless, it can always be used in environments utilising an HTTPS proxy such as data centres or enterprise networks.

IV. EVENTFLOW PROTOTYPE

We build the EventFlow prototype as a plugin for FlowMon [10] flow monitoring software. FlowMon exporter is a flexible flow exporter that provides support for various extensions. These extensions are used to provide support for additional packet inputs, application protocol processing and various export protocols. We utilise the capabilities of the exporter to provide an EventFlow extension plugin.

The FlowMon exporter consists of three main components as shown in Figure 2. The first component is a packet parser. It receives packets from network and extracts information from different layers of each packet. The extracted information is used to create a partial flow record, which contains all necessary information about the parsed packet such as IP addresses, ports, timestamps, byte counter, etc. It also contains application layer information when application parsing is used. The partial flow record is passed to the second component of the exporter which is the flow cache. The partial record is either inserted as a new flow record or it is used to update an existing flow record, which is an aggregation of previous partial flow records. When the flow expires, it is released from the flow cache to the export component. The purpose of the export component is to convert raw flow records to flow export protocol such as NetFlow or IPFIX and pass the flows over network for further processing.

Plugins that extend the FlowMon exporter to process specific application layer protocol such as DNS or HTTP

have access to several parts of the flow creation and exporting process. Each plugin can request to see raw packet payload, process it and add its own information to flow records, such as HTTP Host, Content-Type or Response Code. Furthermore, the plugins are allowed to provide their own functionality for insert, update and release methods of the flow cache. And last, each plugin defines how the information inserted into flow records is exported by the export component.

The EventFlow prototype is implemented in a similar way as an application protocol extension. However, it also utilises the data provided by other application plugins. The EventFlow combines information from DNS and HTTP protocols to detect relations between flows, therefore it requires the DNS and HTTP application plugins to be deployed as well. The prototype extends the packet parser to extract URLs from HTML pages. These URLs are then sent together with partial flow record to the flow cache, however, they are never exported together with the flow record. When new a flow record is created in the flow cache, the expected sets (see Section III) are searched for a match to the new record using DNS, HTTP and URLs information provided in the partial record. When a match is found, the new flow records is assigned an Event ID (8 byte unsigned integer) of the matched record from the expected sets. If no match is found, new EID is incrementally assigned to the new flow record. After the flow is expired from the cache, the EID is a part of the records and is sent by the export component along with the rest of the flow record.

Using 8 byte integer for EID and assigning it incrementally to individual events ensures that there will not be any collisions due to EID overflow in practice. However, an assignment that is individual for each flow probe and persistent over the reboots of the system would be required for a real-world deployment. The EID is assigned only to flows of the HTTP and DNS traffic, since it would provide no benefit to other flows as event relation tracking is not implemented for other protocols yet. Moreover, the size of the flows grows only by 8 bytes at maximum, which has negligible impact on the flow collector disk space requirements.

V. EXPERIMENTAL EVALUATION

We evaluate the prototype in two scenarios. First, we assess the functionality of the prototype on a simple example web site. Once we have verified the functionality, we run EventFlow on a packet trace from live network to determine how many flows can be joined to events in a real traffic. An IPFIXcol [11] flow collector is used to collect and process the generated flows. The collector can be easily configured to work with the Event ID element.

A. Functional Evaluation

For the first scenario, we create a simple website with two pages, each linking the other page, displaying an image and referencing a different JavaScript library. The evaluation proceeds as follows. We request the first page in a browser and few seconds after it loads we follow the link to the other page. The packet trace of these actions is recorded and processed by the EventFlow prototype and the resulting flows are collected by the IPFIXcol.

TABLE II. REAL TRAFFIC EVALUATION STATISTICS.

Total Flows	613953
HTTP Requests	33294
HTTP Responses	49753
DNS Requests	197926
DNS Responses	224588
Events with > 1 Flow	28064
Flows in Events with > 1 Flow	55881
All Events	388749
Flows in All Events	418671

We expect to see a flow record for each of the requests and responses. However, due to HTTP pipelining the whole communication with the web server hosting the test pages is done using single connection. Therefore, there is a pair of flows for the accessing of the two web pages with the linked images (which were on the same server), two pairs of flows flows for each off-site JavaScript library and three pairs of DNS flows for IP address resolution. There are 12 flows created in this test scenario in total. The 12 flows are divided in two user events. The first contains flows for two DNS request, HTTP communication to the web server and download of the first JavaScript library. The second event does not contain an HTTP flow due to the HTTP pipelining but contains DNS request and subsequent download of the second JavaScript library.

The functional evaluation shows that the prototype correctly recognises related flows and allows to mark them as a part of the same event.

B. Real Traffic Evaluation

The purpose of the real traffic test is to determine how many flows can be joined into events. We collect a short (approximately one minute) trace of 10 million packets from an ISP network on ports 53 and 80 which are likely to be DNS and HTTP packets. Table II shows statistics that describe the packet trace as well as the results of the EventFlow prototype. We can see that from the total number of more than 600 thousand flows more than 400 thousand are part of an Event. Furthermore, over 55 thousand flows are part of an event which contains more than one flow. Therefore we can conclude that more than 10 % of observed HTTP and DNS flows are recognised as a part of more complex events by the EventFlow prototype.

The number of flows in complex events is not as high as might be expected since there is a large number of HTTP and DNS requests and responses. We believe that this is caused by quite short time windows of our trace which is likely to have captured only separate responses and requests. Moreover, we believe that better results can be achieved by fine-tuning the timeout of the records in the expected sets of the EventFlow prototype.

VI. CONCLUSIONS

We have presented an EventFlow monitoring architecture that allows to keep track of relations between HTTP and DNS application flows, which can be used to simplify behavioural analysis of network traffic, improve network threat detection and network traffic classification. The changes to existing flow monitoring architecture are negligible, which facilitates wide deployment. The proposed architecture can be further

extended to handle more complex HTTP communication such as redirection return codes.

A prototype of EventFlow plugin for FlowMon flow exporter has been evaluated on a trace of 10 million packets. We showed that more than 10 % of observed HTTP and DNS flows are recognised as a part of more complex events by our prototype. We believe that this result will improve on longer packet trace as well as with more accurate settings of the prototype. Prospective improvements to the prototype as well as its more detailed evaluation are left for a future work.

We believe that the network analysis will benefit from the supplemental information about flow relations. Our work has shown that it is possible to acquire such information without significant impact on existing monitoring architecture and that it is possible extend the flow monitoring to trace the relations of other application protocols.

Acknowledgements

This material is based upon work supported by the “CES-NET Large Infrastructure” project LM2010005 funded by the Ministry of Education, Youth and Sports of the Czech Republic.

REFERENCES

- [1] L. Kekely, J. Kucera, V. Pus, J. Korenek, and A. Vasilakos, “Software Defined Monitoring of Application Protocols,” *Computers, IEEE Transactions on*, vol. PP, no. 99, pp. 1–1, 2015.
- [2] P. Velan, T. Jirsík, and P. Čeleda, “Design and Evaluation of HTTP Protocol Parsers for IPFIX Measurement,” in *Advances in Communication Networking*, T. Bauschert, Ed., vol. 8115. Heidelberg: Springer Berlin Heidelberg, 2013, pp. 136–147.
- [3] ntop, “nProbe,” online, 2015. [Online]. Available: <http://www.ntop.org/products/netflow/nprobe/>
- [4] Y. Wang, Y. Xiang, J. Zhang, W. Zhou, and B. Xie, “Internet traffic clustering with side information,” *Journal of Computer and System Sciences*, vol. 80, no. 5, pp. 1021 – 1036, 2014, special Issue on Dependable and Secure Computing The 9th {IEEE} International Conference on Dependable, Autonomic and Secure Computing.
- [5] H. V. Madhyastha and B. Krishnamurthy, “A Generic Language for Application-specific Flow Sampling,” *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 2, pp. 5–16, Mar. 2008.
- [6] M. Lee, M. Hajjat, R. R. Kompella, and S. G. Rao, “A Flow Measurement Architecture to Preserve Application Structure,” *Comput. Netw.*, vol. 77, no. C, pp. 181–195, Feb. 2015.
- [7] Y. Hu, D.-M. Chiu, and J. C. S. Lui, “Entropy Based Adaptive Flow Aggregation,” *IEEE/ACM Trans. Netw.*, vol. 17, no. 3, pp. 698–711, Jun. 2009.
- [8] L. Dolberg, J. François, and T. Engel, “Efficient Multidimensional Aggregation for Large Scale Monitoring,” in *Proceedings of the 26th International Conference on Large Installation System Administration: Strategies, Tools, and Techniques*, ser. lisa’12. Berkeley, CA, USA: USENIX Association, 2012, pp. 163–180.
- [9] C. Estan, K. Keys, D. Moore, and G. Varghese, “Building a Better NetFlow,” in *Proceedings of the 2004 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, ser. SIGCOMM ’04. New York, NY, USA: ACM, 2004, pp. 245–256.
- [10] INVEA-TECH a.s., “FlowMon Probe,” online, 2015. [Online]. Available: <https://www.invea.com/en/products-and-services/flowmon/flowmon-probes>
- [11] P. Velan and R. Krejčí, “Flow Information Storage Assessment Using IPFIXcol,” in *Dependable Networks and Services*, ser. Lecture Notes in Computer Science, R. Sadre, J. Novotný, P. Čeleda, M. Waldburger, and B. Stiller, Eds., vol. 7279. Heidelberg: Springer Berlin Heidelberg, 2012, pp. 155–158.