Collaborative DDoS Defense using Flow-based Security Event Information

Jessica Steinberger*‡, Benjamin Kuhnert*, Anna Sperotto‡, Harald Baier* and Aiko Pras‡

*da/sec - Biometrics and Internet Security Research Group
University of Applied Sciences Darmstadt, Darmstadt, Germany
Email:{Jessica.Steinberger, Benjamin.Kuhnert, Harald.Baier}@cased.de

‡Design and Analysis of Communication Systems (DACS)
University of Twente, Enschede, The Netherlands
Email:{J.Steinberger, A.Sperotto, A.Pras}@utwente.nl

***Abstract*—Over recent years, network-based attacks evolved to the top concerns responsible for network infrastructure and service outages. To counteract such attacks, an approach is to move mitigation from the target network to the networks of Internet Service Providers (ISP). In addition, exchanging threat information among trusted partners is used to reduce the time needed to detect and respond to large-scale network-based attacks. However, exchanging threat information is currently done on an ad-hoc basis via email or telephone, and there is still no interoperable standard to exchange threat information among trusted partners. To facilitate the exchange of security event information in conjunction with widely adopted monitoring technologies, in particular network flows, we make use of the exchange format FLEX. The goal of this paper is to present a communication process that supports the dissemination of threat information based on FLEX in context of ISPs. We show that this communication process helps organizations to speed up their mitigation and response capabilities without the need to modify the current network infrastructure, and hence make it viable to use for network operators.**

I. Introduction

Nowadays, network-based attacks (e.g. Distributed Denial of Service (DDoS), Distributed Reflection Denial of Service (DrDoS)) pose a serious threat to the network infrastructure and services [1], [2]. To minimize or prevent damages caused by network-based attacks, multiple attack detection methods [3], [4] and countermeasures have been proposed [5], [6]. Recently, these attack detection methods and countermeasures focus more and more on flow data. Besides flow-based attack detection, one approach to counter DDoS/DrDoS attacks focus on collaboration among trusted partners [7], [8]. However, these collaborative approaches do not take into account a communication process that supports the automated exchange of threat information in an interoperable format, uses unreliable and reliable transports and ensures security mechanisms.

In the last years, collaborative approaches have predominantly been published in the area of attack detection [7], but missing to develop collaborative mitigation and response measures. At the same time, several formats (e.g., Incident Object Description Exchange Format (IODEF) [9], Intrusion Detection Message Exchange Format (IDMEF) [10], Abuse Reporting Format (ARF) [11], Extended Abuse Reporting Format (x-arf v0.1 and v0.2) [12] and Flow-based Event eXchange Format (FLEX) [13]) have been published [14] to exchange security events or incidents. However, it is still a challenge to find a standardized exchange format that is thoroughly validated and tested in full scale of industry trails.

978-1-5090-0223-8/16/$31.00 © 2016 IEEE

To overcome the lack of a missing collaborative mitigation and response approach, this paper presents a communication process, that uses the exchange format FLEX [13]. The contribution that the communication process brings to the state of the art is that it supports achieving the situational awareness of the current threat landscape, pools expertise and resources, facilitates the automated defense in response to ongoing network-based attacks and thus lessens the time to respond. In addition, since FLEX messages are disseminated using the Simple (or Streaming) Text Orientated Messaging Protocol (STOMP) or SMTP, the communication process is easy to deploy and integrates with existing infrastructure.

The remainder of this paper is organized as follows: Section II describes the context of our work and its derived requirements. Section III covers the related work by presenting an overview of published collaborative approaches. In Section IV, we introduce the communication process to disseminate security event information among trusted partners, different event producers and consumers, and describe security concerns. In Section V, we analyze and evaluate the communication process. Finally, we conclude the paper in Section VI.

II. Scenario, Requirements and Assumptions

In this section, we describe the main focus of this work. First, we define the networks in which we are going to collaborate among trusted partners and exchange threat information. Second, we define the requirements that a communication process should fulfill, as they emerged by the scenario described in Section II-A. In the following, we will use these requirements to evaluate the communication process. In addition, we describe our assumptions to ensure that the work is not biased by tasks related to detection technologies.

A. Scenario

The primary focus of this work are multiple high-speed networks using a link speed of 10 Gbps and higher [15], and are using an architecture of a typical flow monitoring setup [18] to identify, track and mitigate malicious traffic [16]. The reason to use flow data is, that flow data provides an aggregated view on network traffic passing an observation point [17], and thus reduces the amount of traffic to analyze compared to raw packet data. Further, benefits to use flow-based data are that they are easy to deploy and satisfy the EU

regulation on data retention. In addition, we focus on network operators that cooperate among trusted partners to minimize or prevent damages caused by network-based attacks and use an automated threat information exchange. The main advantage of an automated threat information exchange is to move from a reactive to a proactive network-based attack mitigation and response approach. Another benefit of a collaborative approach is that it provides insight into the current threat landscape that otherwise would not be obvious. In addition, sharing information and collaborating on network-based attacks support to enhance security expertise and speed up the mitigation and response capabilities and thus lessens the time to understand the threat for each collaborating partner.

B. Requirements

In this section, we introduce five requirements that a communication process among trusted partners should fulfill. These requirements are part of the operational requirements described within the IETF DDoS Open Threat Signaling (DOTS) working group [19].

Ease of Deployment: The communication process and its underlaying implementation should support platform independency. Further, the communication process should be able to handle different types of exchange formats and exchange protocols. The platform independency, the use of different exchange formats and protocols ensures that the communication process easily integrates with the existing infrastructure.

Granular access restriction policy: The communication process among trusted partners should support different detail of information tailored for its intended receivers. The reason is that the amount of provided threat information depends on the trust and sharing relationship between collaborating ISPs.

Encryption & **Signature:** The dissemination of security event information often includes sensitive data (e.g., raw data, analyzed information of incident handling and its remediation [20], [21]). Therefore the communication process is required to use an exchange format that supports encryption to prevent unauthorized access to this threat information. Further, the exchange format used within the communication process is required to use signatures to ensure trustworthy origins, relevance and integrity of the security event.

Timeliness: The communication process among trusted partners should ensure the dissemination of security events in an appropriate amount of time. Further, the communication process should move from a reactive to a proactive network-based attack detection and mitigation approach and thus speed up detection and mitigation capabilities.

Semi-automated deployment of countermeasures: The communication process and its underlaying implementation should provide the possibility to interact with the network operator. This interaction ensures that the selection of an automated response depends on the network operator's choice and is called semi-automated. A semi-automated deployment of countermeasures is required to reduce the amount of false automated responses caused by false positives.

C. Assumptions

The detection of malicious actions are performed by monitoring technologies that classify malicious actions based on anomalies [24], [25], [26] or signatures [22], [23]. Both, signature and anomaly-based systems rely on information that determine what is normal behavior and what is not normal. Due to the ever-changing nature of networks, applications, and malicious actions, false positives might be raised. However, the main focus of this work is to show that collaboration among trusted partners helps organizations to speed up their mitigation and response capabilities and thus the assumptions are as follows:

Aggregation: Each alert raised by a detection engine is treated as one attack. The aggregation and fusion of alarms should be taken into account by the detection system. This assumption is in accordance to [27], [28], [29].

Confidence: We assume 100% confidence of the alerts. A detection system might raise false alerts, but to ensure a hundred percent certainty of the alerts or a sanity check of the detection engine is out of scope of this work. This strong confidence is in accordance to [27], [28], [29].

Scalability: We assume that the number of security events raised by the detection engine are not causing a scalability problem, because the quantity of security events that need to be handled by the network operator is low [30].

III. Related work

In this section, we present related work that has been published in the area of collaborative mitigation and response.

a) TAXII and STIX: The Trusted Automated Exchange of Indicator Information (TAXII) is a community-driven effort, that defines concepts, protocols, and message exchanges to share cyber threat information for detection, prevention, and mitigation between trusted partners. The development of TAXII is coordinated by MITRE. The TAXII information exchanged is represented in the XML-based Structured Threat Information Expression (STIX) language [31]. STIX describes potential cyber threat information (e.g., cyber observables, indicators, incidents, adversary tactics, exploits, and courses of action as well as cyber attack campaigns and cyber threat actors) and makes use of Cyber Observable Expression (CybOX) and Common Attack Pattern Enumeration and Classification (CAPEC).

Besides the XML-based structure of STIX, efforts are working on JSON, RDF/OWL, or other implementations. However, STIX and TAXII are complex, require much effort for processing [32] and require extensive adoption to be used within the existing network infrastructure [31]. Further, K. Moriarty [33] reported that the base of TAXII is similar to RID in its use of SOAP-like messaging and thus will likely prevent it from scaling to the demands of the Internet.

b) Firecol: Firecol [7] is a collaborative system that detects flooding DDoS attacks at ISP level and provides a service to which customers can subscribe. This subscription forms a distributed architecture of multiple IPSs that computes and exchanges belief scores on potential attacks. In case Firecol detects an attack, the attack is blocked as close as possible to its source(s). Further, the IPS that detects the attack informs its upstream IPSs, which in turn also performs mitigation procedures. However, Firecol has not been validated

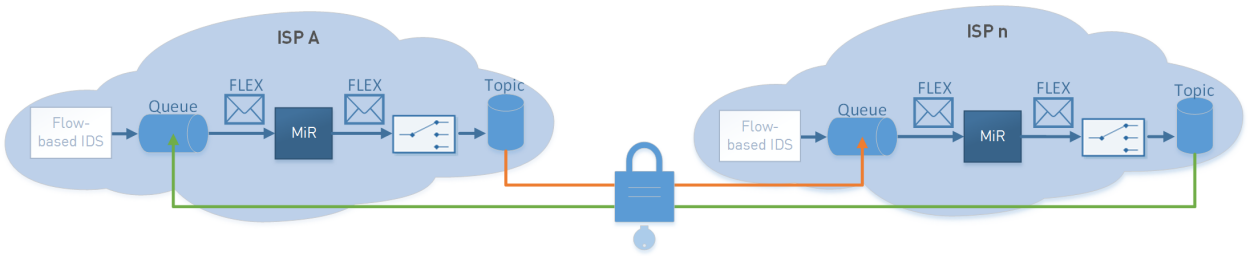

Fig. 1: Exchanging FLEX messages among trusted partners.

in industry trails and its basic theoretic validation is founded on the DARPA'99 dataset [34]. In addition, Firecol only uses blocking as a countermeasure and does not provide additional sophisticated mitigation and response measures.

c) ACDC: In the year 2013, the Advanced Cyber Defence Centre (ACDC) [35] project was launched. ACDC is an effort driven by the European Union to detect, mitigate and respond to botnets. ACDC provides one central database to collect and process data from already existing tools, sensors, sources and further unspecified components. ACDC also supports the mutal data sharing between partners (e.g., ISPs, government agencies, law enforcement, research groups, industry partners). The project finished in June 2015 and provides a web page with documents about the project deliverables. However, the possibility to join the community no longer exists, so the contribution to a collaborative mitigation and response approach remains unclear.

d) Exchange formats and protocols: In recent years, numerous formats (e.g., IDMEF [10], IODEF [9], x-arf [12]) and protocols (e.g., IDXP, BEEP) have been published to support and facilitate the exchange of security event information [14], [31]. Most of the exchange formats are automatically processable to reduce intensive manual processing (e.g, sorting, normalization) and support the timeliness of initial mitigation and response procedures. Moreover, these formats provide an accurate, context rich, directed and actionable data representation of threat information for its intended purpose. Further, as the most formats are based on XML-language or MIME, they ensure integrability with other security tools. Even though network operators often make use of flow-based data to identify, track and mitigate malicious traffic [15], [26], [25], [24], the majority of the published exchange formats are not suitable to convey flow-data without extensive adoption. Besides the ability to convey flow-data, the exchange formats require encryption to prevent unauthorized access to the sensitive data (e.g., raw data, analyzed information of incident handling and its remediation [20], [21]) of the security event. Further, signatures are required to ensure trustworthy origins, relevance and integrity of the security event. The majority of the published exchange formats, except x-arf specification draft v0.2 `X-XARF:SECURE` do not provide any security mechanism.

Another important feature of the exchange format is the ability to support different detail of information tailored for its intended receivers. The reason is that the amount of provided threat information depends on the trust and sharing relationship between two collaborating ISPs.

Despite these formats and protocols are intended to facilitate sharing threat information, it is a challenge to find a standardized exchange format and protocol that is thoroughly validated and tested in full scale of industry trails, because a widespread use of these formats and protocols remain to be established in the community of network operators [14]. A survey performed by [36] and a presentation given by [19] revealed that threat information is often exchanged on an ad-hoc basis via email, member calls or in personal meetings. This slows mitigation and response times and impedes mitigation and reaction efficacy.

IV. Communication Process

In this section, we describe the main components of our proposed communication process and how these components interact with each other.

A. Components of the communication process

Our communication process consists of gateways that are passed by every single security event message. A security event message is transferred among trusted partners using STOMP and the data representation uses the Flow-based Event eXchange format (FLEX). The components of the communication process are illustrated in Figure 1.

a) STOMP: The Simple (or Streaming) Text Orientated Messaging Protocol (STOMP) is a text-based protocol that provides messaging interoperability among many languages, platforms and brokers. Therefore STOMP is language-agnostic and only uses a SEND semantic with a destination string as it does not provide its own queues or topics. STOMP supports messaging features, such as authentication, messaging models (point to point and publish and subscribe), message acknowledgment, transactions, message headers and properties.

The communication process makes use of STOMP between the gateways of the ISPs and is used to transfer FLEX messages among trusted partners. At each destination STOMP is connected to a Java Messaging Service (JMS) queue or topic.

b) FLEX: The Flow-based Event eXchange Format (FLEX) [13] is used to share security information among trusted partners based on flow data. FLEX is based on the x-arf specification draft v0.2 `X-XARF:SECURE` and uses a generic template system that is described by an abstract syntax denoted using the language of Abstract Syntax Notation (ASN.1). In addition, FLEX makes use of both, signature and encryption methods to prevent unauthorized access to the security event message at the application layer.

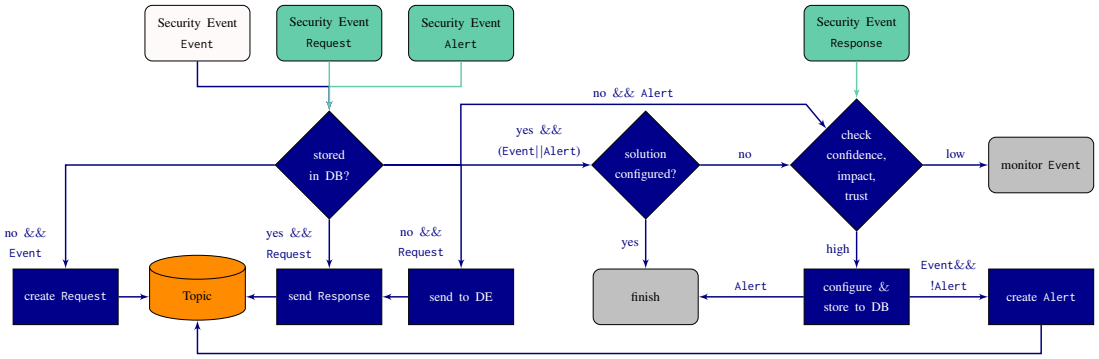

Fig. 2: Data flow of the communication process within MiR from the perspective of one ISP

c) MiR Instances: Each collaborating ISP network contains a mitigation and response system, called MiR. MiR is aligned to the Event Processing Technical Society (EPTS) Reference Architecture [37] and connected to both, a queue and several topics of the JMS. MiR performs pattern matching algorithms to aggregate and consolidate security events into a smaller number of events and thus derives complex events. In addition, MiR enriches the security events through new knowledge gained through previous events or data (e.g., proposes remediations, external publicly available data sources). Besides the queue and the topic, MiR is connected to a database that stores previous security events and their remediation for a defined range of time.

B. Data flow of the communication process

Our communication process consists of interconnected instances located in different ISP networks forming an overlay network. Each ISP possesses a list of directly connected collaborating ISP networks to prevent a full mesh within the network and thus ensure scalability. The data flow of the communication process is shown in Figure 2. The white or green rectangle represents the data entering the communication process (white=internal, green=external), whereas the gray rectangle represents a terminator symbol. The blue rectangle represents a sub-process within MiR. The diamond is used to visualize a decision or branching point and the connected lines represent different options.

At a time t the detection engine of an ISP identifies malicious activities and raises a security event of the FLEX message type `Event`. The security event is sent to the JMS queue of MiR via STOMP. The security event remains in the JMS queue until MiR consumes them. Next, the security event is searched within the database of previous events.

In case the security event could not be found within the database of the previous events, the MiR system initiates a detection process at collaborating neighbor ISP networks to reduce the amount of false positives of the security events. The detection process at the collaborating neighbor ISP networks is initiated by creating a FLEX message of the type `Request` and publishing it to the JMS topic of the adjacent ISP networks. The detection engine of the adjacent ISP network receives the FLEX message of the type `Request`, analyses the network traffic and tries to identify similar behavior as described within the security event. The result of the analysis is sent back to the requesting network as a FLEX message of the type `Response`.

In case the security event could be found within the database of the previous events, the MiR system examines whether the proposed remediation has been configured. In case the proposed remediation has not been configured yet, the ISP evaluates the feedback received in response to the FLEX message `Request` of the connected collaborating ISPs. In case the majority of the connected collaborating ISPs had seen similar behavior, the confidence ranking of the security event is increased, the proposed remediation is configured and stored to the database. Subsequently, the information within the security event is preprocessed and restricted to different level of details depending on the level of trust. Finally, the tailored security events are routed based on their content to the appropriate JMS topic as a FLEX message of type `Alert` and thus send out to the connected collaborating ISPs.

V. Evaluation

In this section, we describe the qualitative and quantitative evaluation of the interaction of the main components of our collaborating MiR system. Further, we evaluate the communication process and its underlying implementation. Finally, we present and summarize the results of the evaluation.

A. Qualitative evaluation

In this section, we perform a qualitative evaluation of the communication process. First, we describe the characteristics of the evaluation criteria. Second, we introduce three evaluation criteria for the communication process.

1) Evaluation methodology: The communication process is evaluated based on the following three criteria: Ease of Deployment, access restriction and security mechanisms (e.g., encryption & signature). These criteria were derived from the requirements described in Section II-B.

The criterion 'Ease of Deployment' describes the ability to use the communication process and its underlying implementation on different operating systems, infrastructure devices, exchange formats and protocols. The criterion 'access restriction' refers to the ability to support different detail of information within an security event tailored for its intended

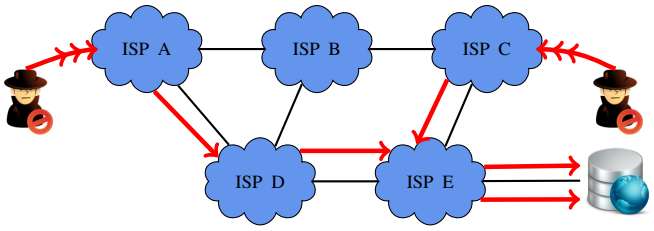

Fig. 3: A semantic representation of the DeterLab testbed

receivers. The 'security mechanisms' criterion describes the ability to make use of an security event exchange format that uses encryption and signature to prevent unauthorized access to the threat information and ensure trustworthy origins, relevance and integrity.

2) Qualitative evaluation results: In this paragraph, we present and discuss the results of the qualitative evaluation of the communication process.

Ease of Deployment: The heterogeneity of network devices and used operating systems requires a platform independent communication process that easily integrates within the existing infrastructure. Therefore the implementation of the communication process is based on Java and thus can easily be deployed on different operating systems. Further, the MiR system was built in a modularized structure and is able to add modules that interact with various network devices. Besides the components of the communication process, the transfer of security events among trusted partners uses STOMP, a language-agnostic protocol to ensure platform independency.

Access restriction: Through the different level of trust and sharing relationship between collaborating ISPs, the communication process supports different detail of information within a security event tailored for its intended receivers. This tailored security event is sent to the appropriate JMS topic to which adjacent ISP networks subscribe and consume security event messages based on their level of trust and sharing relationship.

Security mechanisms: The dissemination of security event information often includes sensitive data (e.g., raw data, analyzed information of incident handling and its remediation [20], [21]). Therefore the communication process uses FLEX to exchange security events among trusted partners, as FLEX supports encryption to prevent unauthorized access to the threat information and uses signatures to ensure trustworthy origins, relevance and integrity of the security event.

B. Quantitative evaluation

In this section, we perform a quantitative evaluation of the communication process. First, we describe the setup of the testbed. Second, we present the test scenario of the communication process.

1) Setup of the testbed: DeterLab [38] is an infrastructure designed for experimentation in context of cyber-security.We used DeterLab to evaluate our collaborative MiR system. The experiment is composed of physical machines for a limited time. DeterLab is used because it is a controlled environment in which it is possible to safely test security threats and defense measures. Our communication process consists of 5 nodes representing Internet Service Providers (ISP $A \dots E$) connected through a 500Mb link. In each ISP network, MiR is installed as shown in Figure 1. Figure 3 shows a schematic representation of the DeterLab testbed.

2) Test scenario: The objective of the experiments is to show that a target network with constrained resources benefits from collaborating partners during an ongoing network-based attack, because the target network has no possibility to react itself due to resource saturation. Further, we show that the network-based attack is not propagated further and that our communication process supports the automatic dissemination of threat information and thus speeds up the mitigation and response capabilities of ISP networks. In addition, we show that our communication process is lightweight in numbers of exchanged messages and fast.

To simulate a network-based attack, we performed a distributed TCP SYN flood attack from the ISP networks A and C to consume resources on the web server within ISP network E and render it unresponsive as shown in Figure 3. In our test scenario, ISP E is not able to effectively block the malicious traffic itself and requires collaborating partners in the stream of traffic to mitigate and respond to the TCP SYN flood attack. To create the TCP SYN flood attack, we used empty TCP packets with a TCP packet size of 40 bytes and a TCP FLAGS value of $0x02$. To ensure that our network-based attack fully utilizes the requested 500Mb link, we sent $40\,000\,000$ TCP SYN packets in total to ensure an attack duration of 26 seconds. However, the internal function of DeterLab does not always allocate resources as requested and thus we received a link connection with 412Mb and were able to perform a TCP SYN flood attack with a duration of 42 seconds.

In our test scenario, the detection engine of the ISP network A initially identifies the malicious network traffic, creates a FLEX message of the type `Event` containing Cisco Netflow version 5 and starts its communication process. The MiR system of the communication process, located in each ISP network, is able to automatically deploy response actions. In our test scenarios, we make use of automatic notifications via email messages including remediation suggestions that make use of `iptables`.In the initial state of the testbed networks no filtering rules are inserted to show the effects of our communication process. During the experiment within DeterLab suspicious IP addresses are identified and inserted into the packet filter ruleset to block the network traffic.

The MiR diagram of the ISP network A in Figure 4 shows four different types of incoming messages. First, ISP A receives two event messages at the same time containing threat information about the network-based attack from ISP network A and C targeting the web server of ISP E. Second, the adjacent collaborating networks report if they had seen similar behavior in their network. Next, ISP A adds two blocking rules to the packet filter ruleset to hit its outgoing links and starts a chain of information that is passed along to ISP C.

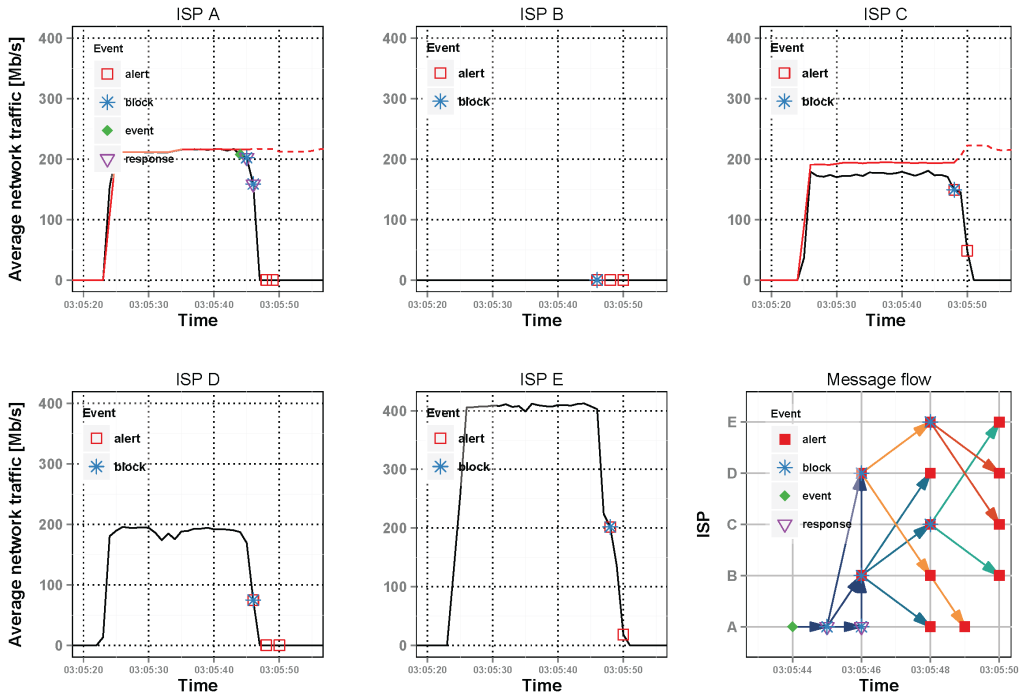

Fig. 4: Simulation results of the attack traffic (red line) and cumulative traffic (black line)

Therefore, ISP A creates an alert and informs the adjacent collaborating networks that deploy the including remediation suggestions. The alerts shown in the MiR diagram of the ISP network A in Figure 4 only represent the incoming alerts and not those that are outgoing. Further, the MiR diagram of the ISP network A in Figure 4 shows that the attack traffic (red line) is sent over 42 seconds. Immediately, after ISP A inserted the blocking rules the effects of the attack traffic are mitigated and the malicious traffic is not propagated further though the network of ISP A (black line). As a result, the traffic at ISP D and E dropped as shown in Figure 4. Next, ISP C deploys the remediation suggestion out of the alert message and as a consequence the traffic at ISP E drops. The MiR diagram of the ISP network E in Figure 4 shows that the collaborating partners in the stream of traffic effectively mitigate and respond to the ongoing network-based attack and thus the network of ISP E is benefiting from the collaboration and the web server recovers.

3) Quantitative evaluation results: In this paragraph, we present and discuss the results of the quantitative evaluation of the communication process.

Timeliness: The primary focus to mitigate and respond to network-based attacks is maintaining the availability of the organization's network infrastructure and services. The Message Flow diagram in Figure 4 shows that the overall duration until all ISP networks have a common knowledge took 6 seconds. Further, 18 messages have been sent until all participating ISP networks had a common knowledge. Even though network B has not been actively involved in the network-based attack, ISP network B has been notified by the ISP network A and thus supports a proactive and collaborative mitigation and response approach. Figure 4 shows that a collaboration among trusted partners facilitates a proactive network-based attack mitigation and response approach and thus contributes to ensure availability of the organization's network infrastructure.

Semi-automated deployment: Through the increasing amount of security events per day, there is a need to automatically process security events and thus lessen the time to mitigate and respond to ongoing network-based attacks. In addition, the automated dissemination of security events among collaborating partners facilitates a proactive network-based attack mitigation and response approach. In our experiment, we have shown that the dissemination of threat information including remediation suggestions exchanged with FLEX are automatically processable and deployable.

VI. Conclusion

Network-based attacks pose a serious threat to the network infrastructure and services. One approach to mitigate and respond to network-based attack focus on collaboration. In this paper, we introduced a communication process that facilitates the automated defense in response to ongoing network-based attacks. We have shown that our communication process is able to proactively mitigate a network-based attack and thus reduces its effects. The main advantage of our communication process over existing approaches is that it easily integrates with the existing infrastructure and is easy to deploy. Based on our qualitative and quantitative evaluation, our communication process constitutes a viable and collaborative approach to disseminate security events among trusted ISP networks. Further, we have shown that our communication process minimizes the complexity of node interactions and will not cause a link congestion.

ACKNOWLEDGMENTS

The work has been funded by the German Federal Ministry of Education and Research (#03FH005PB2), CASED and by EU FP7 Flamingo (ICT-318488).

REFERENCES

[1] Akamai Technologies, Inc., "akamai's [state of the internet] / security Q1 [2015 report]," https://www.stateoftheinternet.com/resources-connectivity-2015-q1-state-of-the-internet-report.html, 2015.

[2] L. Marinos, "ENISA Threat Landscape 2013: Overview of current and emerging cyber-threats," 2013, http://www.enisa.europa.eu/activities/risk-management/evolving-threat-environment/enisa-threat-landscape-2013-overview-of-current-and-emerging-cyber-threats/at_download/fullReport.

[3] V. Sekar, N. Duffield, O. Spatscheck, J. van der Merwe, and H. Zhang, "LADS: Large-scale Automated DDoS Detection System," in *Proceedings of the Annual Conference on USENIX '06*, 2006.

[4] G. Münz and G. Carle, "Real-time Analysis of Flow Data for Network Attack Detection," in *Proceedings of the 10th IFIP/IEEE International Symposium on Integrated Network Management (IM 2007)*, May 2007.

[5] H. Beitollahi and G. Deconinck, "Analyzing well-known countermeasures against distributed denial of service attacks," *Computer Communications*, no. 11, 2012.

[6] C. Rossow, "Amplification Hell: Revisiting Network Protocols for DDoS Abuse," in *Proceedings of the 2014 Network and Distributed System Security (NDSS) Symposium*, February 2014.

[7] J. François, I. Aib, and R. Boutaba, "Firecol: A collaborative protection network for the detection of flooding DDoS attacks," *IEEE/ACM Transactions on Networking*, no. 6, Dec 2012.

[8] G. Oikonomou, J. Mirkovic, P. Reiher, and M. Robinson, "A Framework for a Collaborative DDoS Defense," in *Proceedings of the 22nd Annual Computer Security Applications Conference*, 2006.

[9] R. Danyliw, J. Meijer, and Y. Demchenko, "The Incident Object Description Exchange Format RFC 5070 (Proposed Standard)," IETF, Dec. 2007.

[10] H. Debar, D. Curry, and B. Feinstein, "The Intrusion Detection Message Exchange Format (IDMEF) RFC 4765 (Experimental)," IETF, Mar. 2007.

[11] Y. Shafranovich, J. Levine, and M. Kucherawy, "An Extensible Format for Email Feedback Reports RFC 5965 (Proposed Standard)," IETF, Aug. 2010.

[12] abusix GmbH, "x-arf Network Abuse Reporting 2.0," http://www.x-arf.org/.

[13] J. Steinberger, A. Sperotto, H. Baier, and A. Pras, "Exchanging Security Events of flow-based Intrusion Detection Systems at Internet Scale," Internet Architecture Board and the Internet Society, https://www.iab.org/wp-content/IAB-uploads/2015/04/CARIS_2015_submission_3.pdf , June 2015.

[14] J. Steinberger, A. Sperotto, M. Golling, and H. Baier, "How to Exchange Security Events? Overview and Evaluation of Formats and Protocols," in *2015 IFIP/IEEE International Symposium on Integrated Network Management (IM 2015)*, May 2015.

[15] M. Golling, R. Hofstede, and R. Koch, "Towards multi-layered intrusion detection in high-speed networks," in *6th International Conference On Cyber Conflict, 2014*, June 2014.

[16] J. Steinberger, L. Schehlmann, S. Abt, and H. Baier, "Anomaly Detection and Mitigation at Internet Scale: A Survey," in *Emerging Management Mechanisms for the Future Internet*, 2013.

[17] A. Sperotto, G. Schaffrath, R. Sadre, C. Morariu, A. Pras, and B. Stiller, "An Overview of IP Flow-Based Intrusion Detection," *IEEE Communications Surveys Tutorials*, vol. 12, no. 3, 2010.

[18] R. Hofstede, P. Celeda, B. Trammell, I. Drago, R. Sadre, A. Sperotto, and A. Pras, "Flow Monitoring Explained: From Packet Capture to Data Analysis with NetFlow and IPFIX," *IEEE Communications Surveys & Tutorials*, 2014.

[19] C. Morrow and R. Dobbins, "DDoS Open Threat Signaling (DOTS) Working Group Operational Requirements," IETF 93, https://www.ietf.org/proceedings/93/slides/slides-93-dots-3.pdf, July 2015.

[20] R. Ruefle, A. Dorofee, D. Mundie, A. Householder, M. Murray, and S. Perl, "Computer Security Incident Response Team Development and Evolution," *IEEE Security & Privacy*, no. 5, Sept 2014.

[21] K. Moriarty, "Incident coordination," *IEEE Security & Privacy*, no. 6, Nov 2011.

[22] C. Sanders and J. Smith, *Applied Network Security Monitoring - Collection, Detection, and Analysis*, 1st ed. Syngress: Elsevier, December 2013.

[23] M. A. Faysel and S. S. Haque, "Towards cyber defense: research in intrusion detection and intrusion prevention systems," *IJCSNS International Journal of Computer Science and Network Security*, vol. 10, no. 7, 2010.

[24] J. François, S. Wang, R. State, and T. Engel, "BotTrack: Tracking Botnets Using NetFlow and PageRank," in *NETWORKING 2011*, 2011.

[25] F. Tegeler, X. Fu, G. Vigna, and C. Kruegel, "BotFinder: Finding Bots in Network Traffic Without Deep Packet Inspection," in *Proceedings of the 8th International Conference on Emerging Networking Experiments and Technologies*, 2012.

[26] L. Bilge, D. Balzarotti, W. Robertson, E. Kirda, and C. Kruegel, "Disclosure: Detecting Botnet Command and Control Servers Through Large-scale NetFlow Analysis," in *Proceedings of the 28th Annual Computer Security Applications Conference*, 2012.

[27] C. Strasburg, "A framework for cost-sensitive automated selection of intrusion response," Master's thesis, Iowa State University, 2009.

[28] C. Strasburg *et al.*, *A Framework for Cost Sensitive Assessment of Intrusion Response Selection*. Institute of Electrical and Electronics Engineers, July 2009.

[29] F. P. Stanley, "Intrusion detection and response for system and network attacks," Master's thesis, Iowa State University, 2009. [Online]. Available: http://lib.dr.iastate.edu/cgi/viewcontent.cgi?article=1730&context=etd

[30] J. Steinberger, A. Sperotto, H. Baier, and A. Pras, "Collaborative Attack Mitigation and Response: A survey," in *2015 IFIP/IEEE International Symposium on Integrated Network Management (IM 2015)*, May 2015.

[31] P. Kampanakis, "Security automation and threat information-sharing options," *IEEE Security & Privacy*, no. 5, Sept 2014.

[32] DFN-CERT Services GmbH, "D1.7.1 – Data Format Specification," http://acdc-project.eu/wp-content/uploads/2015/05/ACDC_D1.7.1_Data_Format.pdf , July 2013.

[33] K. M. Moriarty, "CARIS Workshop Summary and Reflection," Internet Architecture Board and the Internet Society, https://www.ietf.org/blog/2015/06/caris-workshop-summary-and-reflection/ , June 2015.

[34] R. Lippmann, J. W. Haines, D. J. Fried, J. Korba, and K. Das, "The 1999 DARPA Off-Line Intrusion Detection Evaluation," *Comput. Netw.*, no. 4, 2000.

[35] Advanced Cyber Defence Centre, "ACDC Deliverables," http://acdc-project.eu/acdc-deliverables/ , 2015.

[36] Internet Architecture Board and the Internet Society, "CARIS Workshop Template Submissions," Internet Architecture Board and the Internet Society, https://internetsociety2.wufoo.com/reports/caris-workshop-template-submissions/, June 2015.

[37] A. Paschke and P. Vincent, "A reference architecture for event processing," in *Proceedings of the Third ACM International Conference on Distributed Event-Based Systems*, 2009.

[38] J. Mirkovic and T. Benzel, "Teaching Cybersecurity with DeterLab," *IEEE Security & Privacy*, vol. 10, Jan 2012.