

Conceptual Integration of Flow-based and Packet-based Network Intrusion Detection

Gregor Schaffrath, Burkhard Stiller

Department of Informatics IFI, University of Zürich
Communication Systems Group CSG
Binzmühlestrasse 14, CH—8050, Zürich, Switzerland
schaffrath@ifi.uzh.ch, stiller@ifi.uzh.ch

Abstract. Network-based Intrusion Detection Systems aim at the detection of malicious activities by an inspection of network traffic. Since network link speeds and traffic volume grew over the last years, payload-based analysis became difficult, leading to the development of alternative approaches for flow-based analysis. Although each approach alone suffers a set of drawbacks, a few experiments with hybrid approaches show potential for synergies. This work analyses these drawbacks in order to develop a conceptual framework for hybrid approaches, integrating the two concepts in a fashion to compensate for their respective weaknesses proposed.

1 Introduction, Motivation, and Goals

Network-based Intrusion Detection Systems (NIDS) aim at the detection of malicious activities by inspection of network traffic. As network link speeds and traffic volume continued to grow, the traditional approach of Packet-based (or Payload-based) NIDS (PNIDS) became difficult [2]. Alternate Flow-based NIDS (FNIDS) approaches were developed, but while the development of evaluation methods are still ongoing [1], they are generally recognized as featuring a reduced confidence level expressed by higher false positive rates and do not allow for a fine grained labelling, as PNIDS concepts do.

Although few experiments with hybrid approaches indicate synergy potential [4], PNIDS and FNIDS seem to be rarely considered in combination in current research literature and seem to be perceived rather as competing than complementary concepts.

The goal of this work is the development of a conceptual framework for hybrid approaches. Strengths and weaknesses of both PNIDS and FNIDS are analyzed. Sample scenarios are developed, in which characteristics of PNIDS and FNIDS may be mapped onto requirements in order to compensate for respective drawbacks. These are generalized into abstract use templates based on a conceptual understanding of common issues and requirements, enabling easier design of security concepts by instantiation rather than complete redesign for each scenario.

2 Analysis

In an ideal world without any resource constraints, the set of attacks detectable by FNIDS, both on suspicion as well as on certainty level, are a full subset of attacks detectable by PNIDS. This is the consequence of FNIDS operating on a subset of infor-

mation available to PNIDS. However, in real world environments, PNIDS is confronted with several problems reducing its effectiveness, resulting in a separation the of sets of detectable attacks for both approaches, as presented in Figure 1. These problems can be summarized in two categories:

- **Resources:** Packet-based inspection in high volume networks requires many resources in terms of memory and usually also CPU time. This is due to the low level of abstraction of the data basis in combination with the high frequency of data arrival and possible large data quantities per information instance.
- **Availability:** The placement of monitoring devices is crucial. However, if analysis requires transfer of monitored packets to a remote location, availability easily becomes an issue to several respects. Examples are bandwidth constraints for these transfers or privacy concerns on the transmission of actual payload information from production systems.

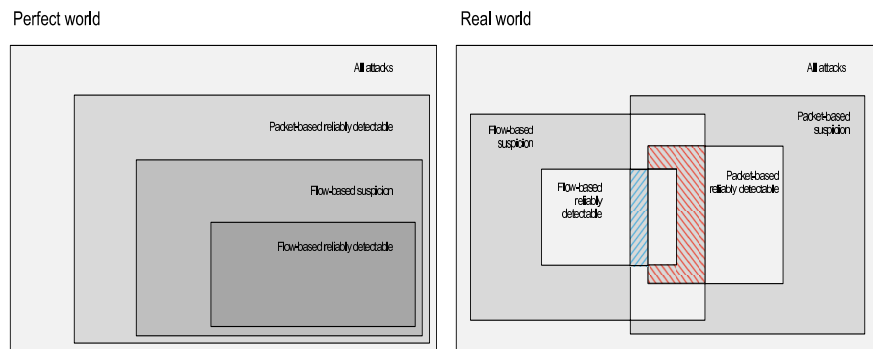


Figure 1 Detection Coverage [3]

While the resource issue has been recognized by the research community, attempts to solve it (*e.g.*, [6]) have been largely focused on the OS and tool level. No work addressing it by a conceptual combination with other approaches has been established on the analysis level.

FNIDS is in a better position to handle large traffic volumes, since it processes a smaller amount of information and allows for the delegation of assembly work of input data even to regular network devices like switches. These two characteristics also increase their potential w.r.t. data availability, since device installation is less of a problem, data transfer is easier to handle and privacy concerns are less important. Nonetheless, it suffers a set of problems as well, which can be summarized by the following three categories:

- **Confidence:** In a complex context like security, where every arbitrary detail may be crucial to accurate analysis, the reduced amount of information can be expected to result in the experienced drop of confidence (resulting in higher false positive rates) or alert expressiveness (*e.g.*, allowing statements only about activity categories, instead of concise attack labelling)

- **Underdeveloped understanding:** Attack characteristics on the flow level are not yet thoroughly analyzed. This is also reflected by the fact that available research documentation shows a trend to visualization concepts (*e.g.*, [5]) and almost all work in FNIDS remains anomaly-based. Misuse model-based concepts remain underdeveloped and are the exception in FNIDS research.
- **Real-time:** When information assembly is delegated to the network infrastructure delivering, *e.g.*, NetFlow records in order to take full advantage of the resource requirement reduction advantage, reports will be delayed after flow's end w.r.t. reported activities, severely reducing FNIDS usability for real-time intrusion detection.

3 Proposed Combination Angles

Since flow information reflects a high level view on the interaction behavior of network nodes with each other, but seldomly allow for statements about specific instances of these interactions, they can be deemed predestined for high level characterizations of host behavior or roles, while packet information can be expected to be fit for directed in-depth investigations.

Considering distribution issues and classifying along resource and availability issues mentioned in the PNIDS analysis yields to two basic combination angles, where synergy effects may be expected:

- Multi-stage concepts for resource efficiency and
- Coordination concepts in distributed environments.

Multi-stage concepts for resource efficiency use results of one approach as the input for directed investigations on the basis of the other approach. Depending on the network environment and specific protection goals determined in advance, this can be beneficial in both directions:

FNIDS may be used as a selection filter for PNIDS activities for CPU or memory requirement reasons. Example use cases (to be evaluated) for this include the isolation of P2P traffic, recognized by flow level characteristics, for packet level inspection of suspicious flows, in order to tell regular P2P traffic apart from P2P botnet traffic, or the use of PNIDS for labelling purposes upon FNIDS based worm recognition. Depending on scenario requirements and available capacities to capture and cache all relevant packet data, the follow-up inspection may either be restricted to the subsequent communication of hosts involved or be performed on the actual data that triggered the flow-based alert.

Critical assets may be protected by PNIDS for real-time reasons, whose alert confidence upon a signature trigger could be increased by judgement on the overall behaviour of hosts involved. Example use cases for this include (a) the search for repetitive host behavior upon a alert triggered from a generic payload signature indicating worm activities, or (b) the flow-based search for control streams upon botnet detection, where flows of hosts involved are checked for common communication partners and properties. This is exemplified in [4], while the analysis encompasses possibly the communication between different administrative zones and additional subsequent PNIDS.

In some scenarios, coordination concepts in distributed environments are related to multi-stage concepts. This may be based on network resource-related problems, where efficiency depends on which information is analyzed locally w.r.t. each site and which information is forwarded to another site. An example of this is a scenario, in which a customer is connected to the Internet via two different Internet service providers, for load balancing or redundancy reasons, and parts of attacks are distributed over uplinks. In this case, flow information could be exchanged in order to correlate flows on both links and start packet transfers for PNIDS-based inspection of suspicious traffic. *E.g.*, upon the detection of flows with equal endpoints, these flows might be checked locally for packet-level fragmentation, followed by a potential packet data exchange for reassembly.

However, in many scenarios, coordination concepts involve additional concerns and issues not covered by scenarios for multi-stage concepts: While the transfer of information may not be a technical issue, payload transfer from one site to another may be impossible and even the transfer of flow-based information may be restricted, *e.g.*, by policies. This shift of focus w.r.t. concerns can be expected to result in shifted use concepts of PNIDS and FNIDS. An example for this is a scenario, where an administrative domain A registers an anomaly triggered by a host in domain B, and suspects a correlation to the activity registered from domain C. Since the domains are prohibited to exchange payload information, they are restricted to tentative correlation via FNIDS or delegation of PNIDS tasks to individual sites.

4 Evaluation

While evaluation of the effectiveness of the hybrid approaches may be relatively straightforward for experimentation w.r.t. resource efficiency, at the time of writing, it is yet an open point, how to evaluate the quality of distributed scenarios and their respective templates and solutions.

Resource efficiency of multi-stage concepts may be determined by running both a simple packet-based IDS and the combined approach simultaneously in the same environment on equal infrastructure and measuring respective resource consumption, as well as alert, false positive and false negative rates. If the processed traffic is reduced to an amount manageable by both approaches, the direct correlation of overall resource consumption with achieved results concerning alert correctness may serve as quality measure.

However, while quality in distributed environments may be measurable by similar metrics to some extent, several aspects are yet unaddressed and, therefore, still open in terms of the evaluation. *E.g.*, the scenario choice, especially in terms of environments, detection goals, and attack scopes, needs to be validated for realism. Solutions need to be validated w.r.t. policy compatibility (and, therefore, applicability), and real world evaluation will depend on the availability of cooperation partners.

As these scenarios are potentially arbitrarily numerous and complex, it is intended to start by isolating motivations for distributed Intrusion Detection and mapping them to real world instances of distributed systems. Thereby, a break down of the problem into a manageable complexity will follow a practical track of solutions.

5 Preliminary Conclusions, Tasks, and Issues

Following the two combination angles, the work will be split into two parts. The investigation into the resource effectiveness of multi-stage approaches consists of:

- the definition of use scenarios for comparative evaluation of resource and detection efficiency in a real world test environments,
- the generalization into use templates, if results confirm the efficiency hypothesis,
- re-instantiation of templates into different use scenarios for validation of the generalization.

Investigations of coordination concepts in distributed environments consist of:

- generation of realistic prototypical cooperation scenarios,
- development of approaches to handle political, legal, and operational issues,
- integration of concepts into use templates based on respective scenarios

Investigations and evaluations w.r.t. multi-stage approaches for resource efficiency are currently ongoing in the context of [3].

Acknowledgements. The work of Fabian Hensel in support of first steps for developing the evaluation scenarios and implementation environments is acknowledged.

References

- [1] D. Brauckhoff, M. May, B. Plattner: *Flow-Level Anomaly Detection - Blessing or Curse?*, IEEE INFOCOM 2007, Student Workshop, Anchorage, Alaska, U.S.A., May 2007.
- [2] H. Dreger, A. Feldmann, V. Paxson, R. Sommer: *Operational experiences with high-volume network intrusion detection*, 11th ACM Conference on Computer and Communications Security, Washington, U.S.A., 2001.
- [3] F. Hensel: *Flow-based and Packet level-based Intrusion Detection as Complementary Concepts*, Diploma Thesis, University of Zurich, Department of Informatics IFI, Switzerland, April 2008.
- [4] A. Karasaridis, B. Rexroad, D. Hoeflin: *Wide-scale Botnet Detection and Characterization*, HotBots '07, Usenix Workshop on Hot Topics in Understanding Botnets, Cambridge, Massachusetts, U.S.A., April 2007.
- [5] K. Lakkaraju, W. Yurcik, A. J. Lee: *NVisionIP: netflow visualizations of system state for security situational awareness*, 2004 ACM Workshop on Visualization and Data Mining for Computer Security, Washington D.C., U.S.A., 2004.
- [6] F. Schneider: *Performance evaluation of packet capturing systems for high-speed networks*, Diploma Thesis, Technische Universität München, Munich, Germany, November 2005.