Visual Analytics of Network Security Metadata

Volker Ahlers Bastian Hellmann
University of Applied Sciences and Arts Hannover
Germany

Abstract – Many computer network components provide security-relevant metadata, ranging from log files to the output of intrusion detection systems. Different visualization approaches exist, which mainly concentrate on single components. Based on the IF-MAP protocol, which defines an integrated, graph-based view on computer networks, we introduce a visual representation of the whole network dynamics. We present a visual analytics framework that allows to determine which network components are related to a security-critical event. Detection policies can be checked by means of a "what-if" simulation, showing the effects of configuration changes on the analysis of past network dynamics and thus facilitating the minimization of false positive detection results. The presentation gives an overview of our own work of recent years as well as related work, closing with a glimpse on future research directions.

1. Introduction and Related Work

Analyzing the security of computer networks is a challenge due to their dynamic nature. Users are logging on and off, mobile devices are temporarily attached, and applications and services are updated at irregular intervals. Usually several different network security detection systems such as firewalls, network access control (NAC) components, and intrusion detection systems (IDS) exist in parallel to monitor certain aspects of network security. Some attacks, however, are characterized by a combination of these aspects and can thus only be detected by an integrated view.

In recent years several visual analytics approaches to network security monitoring have been proposed, such as visualizing relations between hosts, users, and applications [1], visualizing attack graphs of detected events to foster visual awareness [2], and security dashboards visualizing information from different sources [3].

What is missing in these approaches, however, is an integrated view on the available data in a consistent visual abstraction. In the following we describe an open-source framework using an integrated, extensible data model to this end. Furthermore, we present an approach to adjust security policies based on a historical view on the network data and results of a user study. This paper gives an overview of our framework. Implementation details can be found in the references.

2. Visual Analytics Framework

The fundamental approach of our visual analytics framework VisITMeta is to collect security-relevant network data from different sensors and to store this data in an integrated data model for further analysis. Figure 1 depicts the general framework architecture. The data model is based on the Interface for Metadata Access Points (IF-MAP)

specification [4]. Within this specification a MAP server collects data from various physical and logical network components (MAP clients), using a graph-based data model consisting of two types of nodes – identifiers and metadata – which are connected by edges. Identifiers represent physical and logical network entities such as devices, IP and MAC addresses, or users. Metadata specify connections between identifiers by, e.g., user roles or IP-MAC associations. Furthermore, metadata are used to describe detected events such as network attacks or security threats. The framework is implemented in Java and available as open-source software via GitHub [5].

Since IF-MAP data are distributed with a common timestamp, it is possible to compare network states at different time instances. Our framework offers a delta view, which marks network entities added or removed during a selected time period by means of colored glow effects, as shown in figure 2. This enables the network administrator to detect possible security threats in concordance with suspicious network changes, or to compare the normal network state to a changed state after a detected attack in order to find the source of the attack [6].

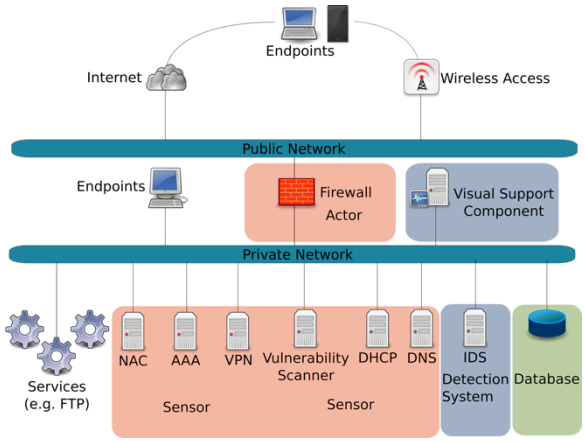

Figure 1: Framework architecture consisting of sensors and actors (red), analysis components (blue), and data storage (green).

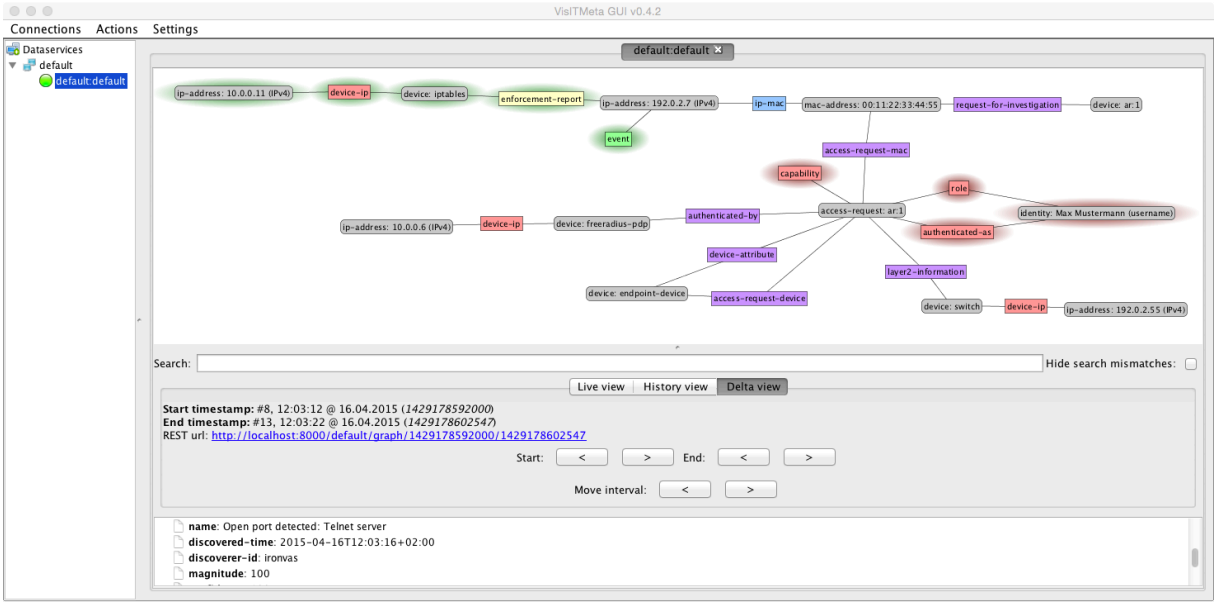

Figure 2: The VisITMeta user interface with an example graph visualized in force-directed layout. Identifiers are shown as nodes with round edges, metadata as nodes with sharp edges. In delta view mode, green and red glow effects mark nodes that have been added and removed during a certain time period, respectively. In the bottom section of the interface, details for a selected node (the green event) are shown.

3. "What-if" Policy Adjustment

The VisITMeta framework offers the visualization of a detected security event together with the current network state and the detection policy that triggered the event, thus enabling the network administrator or security specialist to trace the source and evaluate the plausibility of a detected event. For consistency, the detection policies consisting of patterns and rules are modeled as IF-MAP graphs, as shown in figure 3. In order to reduce false-positive events or care for known false-negative, i.e., undetected events, the policy rules can be adjusted and tested with historical network data via a "what-if" simulation using data exchange mechanisms similar to the IF-MAP specification [7].

4. User Study Evaluation

To evaluate our framework, we carried out a user study with 12 test subjects with different background knowledge (undergraduate, graduate, and Ph.D. students, network administrators). All subjects were given a list of tasks to be carried out in three different example scenarios, e.g., finding out details of the network state, checking rules of detection policies, or adjusting a policy configuration. Some of the tasks had to be carried out first based on isolated information such as log data and then with the VisITMeta framework in comparison.

After completing the tasks, the subjects filled out a questionnaire partially based on the system usability scale [8]. Furthermore, the success as well as the required time for the tasks were analyzed. In summary, most of the tasks were completed correctly (63.7%) or partially correctly (30.4%), while only a small number was completed incorrectly (3.0%) or left unanswered (3.0%). In the answers to the questionnaire, most subjects agreed that the integrated view of detected events and detection policies is useful, with average ratings between 2.6 and 3.3 on a scale from 0 (no agreement) to 4 (full agreement) [9].

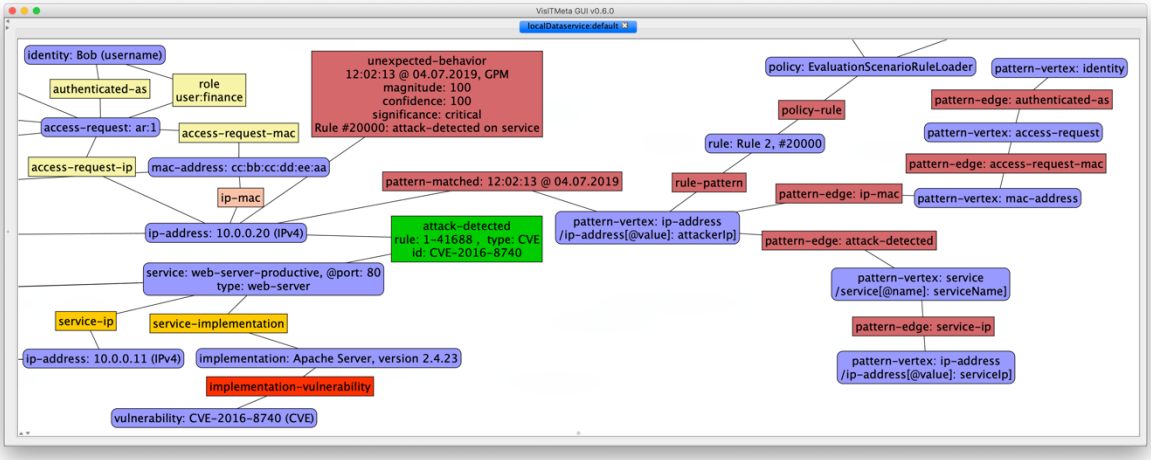

Figure 3: Example of VisITMeta visualizing a detected security event (green) together with the current network state (left) and the policy that triggered a security event (right), which are connected by a "pattern-matched" metadata node (center).

5. Conclusion

We have presented a visual analytics framework for network security information from different sources, using an integrated data model based on the IF-MAP specification. The open-source framework VisITMeta [5] allows the comparison of network states at different time instances as well as the adjustment of detection policies based on "what-if" simulations with historical data.

An open task is the scalability of the approach for large networks. This could be tackled by a level of detail mechanism that hides irrelevant information. A further direction of current work is the application of machine learning techniques for the creation of detection rules.

Acknowledgment

This work was financially supported by the German Federal Ministry of Education and Research (BMBF), projects VisITMeta (grant no. 17PNT032) and SIMU (grant no. 16KIS0045). The fruitful collaboration with Felix Heine, Carsten Kleiner, the members of the Trust@HsH group, and our industry partners is gratefully acknowledged.

References

[1] Q. Liao, A. Striegel, and N. Chawla, "Visualizing graph dynamics and similarity for enterprise network security and management," in Proc. VizSec. ACM, 2010, pp. 34-45.

[2] M. Angelini, N. Prigent, and G. Santucci, "Percival: proactive and reactive attack and response assessment for cyber incidents using visual analytics," in Proc. VizSec. IEEE, 2015, pp. 1–8.

[3] J. R. Goodall, E. D. Ragan, C. A. Steed, J. W. Reed, G. D. Richardson, K. M. T. Huffer, R. A. Bridges, and J. A. Laska, "Situ: Identifying and explaining suspicious behavior in networks," IEEE Transactions on Visualization and Computer Graphics, vol. 25, no. 1, 2019, pp. 204–214.

[4] Trusted Network Communications Working Group, "TNC IF-MAP binding for SOAP, version 2.2, revision 10," Trusted Computing Group, March 2014.

[5] Trust@HsH Group, "Iron/VisITMeta projects on GitHub," https://github.com/trustathsh/, last accessed 2019-09-26.

[6] V. Ahlers, F. Heine, B. Hellmann, C. Kleiner, L. Renners, T. Rossow, and R. Steuerwald, "Integrated Visualization of Network Security Metadata from Heterogeneous Data Sources," in Proc. GraMSec. Springer International Publishing, 2016, pp. 18–34.

[7] B. Hellmann, V. Ahlers, and G. Dreo Rodosek, "Integrating visual analysis of network security and management of detection system configurations," in Proc. IDAACS. IEEE, 2017, pp. 1020–1025.

[8] J. Brooke, "SUS–a quick and dirty usability scale," Usability Evaluation in Industry, vol. 189, no. 194, 1996, pp. 4–7.

[9] V. Ahlers, B. Hellmann, and G. Dreo Rodosek, "A user study of the visualization-assisted evaluation and management of network security detection events and policies," in Proc. IDAACS. IEEE, 2019, pp. 668–673.

