Toward Effective Visualization of Network Identifier Bindings in a Software-Defined Network

Matt Gottsacker* Steven R. Gomez# Rick Skowyra# Flavio Esposito*

*Saint Louis University, USA #MIT Lincoln Laboratory, USA

ACM Internet Measurement Conference 2018

Abstract

The dynamic nature of network-level bindings, for example among usernames, hostnames, IP and MAC addresses, complicates already difficult network analysis tasks. One such task is determining which users are bound to which IP addresses in order to verify an access control rule in a Software-Defined Network (SDN). Existing SDN verification tools merely display the topology of network devices [1, 2]. As a consequence, it is typical for network administrators to gather additional identifier information by examining logs from services or sensors. Because of the scale of an enterprise network, this method is inefficient and may lead to change blindness. We present work toward a novel visualization tool to combat these problems.

In particular, we approached these issues by designing a new web application that displays two different visual representations of the identifier binding database. This database contains tables of pairs of bound identifiers. While this identifier binding database is not implemented in typical networks, it is reasonable for network administrators to maintain such a structure to perform better forensics on the network following a security event. Each visualization method shows an overview of the network to allow top-down analysis, as well as features that give users access to the granular data necessary for bottom-up analysis.

Identifier Mapping Construction

Tracing mappings among network identifiers requires knowledge of several different binding relationships. Different network services and sensors discover each pair of bindings. A database is useful for keeping track of the accumulative mappings.

Switch ID	MAC
23525	AB:31:...
23457	1F:34:...

Packet Inspection

MAC	IP
AB:31:...	122.45.
1F:34:...	154.120.

DHCP

DNS

IP	Hostname
122.45.	foobar
154.120.	bazqux

Logon Sensor

Hostname	Username
foobar	alice
bazqux	bob

Method

We designed and implemented a web-based application with two interactive visualizations to simplify network analysis tasks [3]. Each one encodes metadata about the validity of each binding, e.g., expiration timestamps as colors along a colorblind-safe gradient. The freshest bindings are purple, while stale bindings become orange. When a binding is updated, the identifiers associated with it are highlighted. Additionally, the visualizations have zooming options to help them scale with enterprise networks. These features give analysts quick insights about the state of the network and help them observe changes to it.

Table Approach: Intuitively Sort and Filter Identifiers

The first representation is an enriched tabular view, in which records from a database of current identifier bindings are fit into a sortable and filterable table. Identifier names are columns that have been joined from database tables representing individual binding pairs. The link icons between identifier columns are colored along the gradient described above. This view is intuitive because each row represents a distinct entity that is known to the network; however, because multiple entities might share an identifier, the table length could be prohibitively large for an analyst.

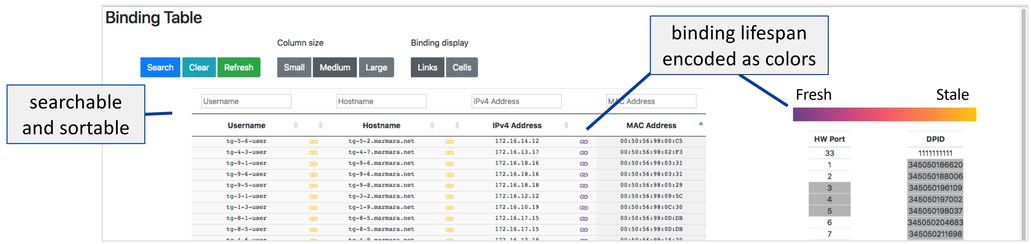

Graph Approach: Visualize Interesting Questions

The second representation displays an undirected graph with the unique network identifiers as nodes and the bindings as edges between the corresponding identifiers. This view is intended to support analyst questions that are focused on individual identifiers. While this view loses some simplicity, it avoids the scaling problem of showing a unique identifier in multiple places, and the analyst can learn at a glance whether a certain identifier is bound to several others.

Conclusion

We have implemented and begun testing a web-based application for the views described above. The tool updates in real-time to ensure the network analyst has up-to-date information. The next steps for this work include evaluating the efficacy of the system with real analysts in order to better understand the benefit of the tool to the network measurement community.

Applications

Each of these visualization techniques is useful for answering questions that network analysts may have about the state of the network. Analysts often need to understand the hosts or users behind network traffic. Our visualization tool provides this information faster and in a more meaningful way than traditional verification methods.

Future Work and Open Challenges

- User studies to determine the effectiveness of the visualizations and other visualization techniques.
- Visualizing how long bindings have been expired and limiting the views based on de/activated timestamps. These features could be useful when analysts conduct provenance investigations after a security event.
- Combine the binding visualizations with views that visualize policies to gain more insights into network reachability at any given moment.

References

[1] ”Floodlight Web GUI”, May 6, 2016, https://floodlight.atlassian.net/wiki/spaces/floodlightcontroller/overview/.

[2] “The ONOS Web GUI”, May 17, 2017, https://wiki.onosproject.org/display/ONOS/.

[3] D. Staheli, T. Yu, R. J. Crouser, S. Damodaran, K. Nam, D. O'Gwynn, S. McKenna, L. Harrison, "Visualization evaluation for cyber security: Trends and future directions", *Proceedings of the Eleventh Workshop on Visualization for Cyber Security*, pp. 49-56, 2014.

Acknowledgements

This work has been partially supported by the NSF award CNS-1647084. M. Gottsacker is an undergraduate student supported by an REU.

This work was completed while M. Gottsacker was at MIT Lincoln Laboratory.