Metadata driven reconstruction of enterprise network segmentation from communication logs

Philipp Neumann
Faculty of Mathematics and Computer Science
FernUniversität in Hagen
Hagen, Germany
0009-0004-3038-1876

Prof. Dr. Jörg Keller
Faculty of Mathematics and Computer Science
FernUniversität in Hagen
Hagen, Germany
0000-0003-0303-6140

***Abstract*—Internal networks of large enterprises are segmented for security reasons, but centralized information about segment structure is often incomplete or outdated as networks evolve over time. This might lead to security gaps and prevents optimal setup of segments in current enterprise networks. We propose to approximate and/or complete the segment structure via log data available from routers, firewalls and other network devices, thus getting more accurate information that helps to identify security gaps and serves as a starting point for further network structure optimization. As a first step, we propose an algorithm to reconstruct an enterprise's segment structure as a list of security zones from communication logs on firewalls in the enterprise network. We evaluate the algorithm both with synthetic and real network data, for which the segment structure of two large companies has been constructed manually, i.e., where the ground truth is known. Our results demonstrate that already a restricted set of communication logs suffices to reconstruct a large part of the network structure and thus provides a fast and automatic way to update and complete network documentation so that further optimization with known algorithms is possible.**

***Index Terms*—network segmentation, network architectures, network modeling**

I. Introduction

Network segmentation is widely recognized as a fundamental security mechanism in modern enterprise IT environments. By limiting unrestricted communication to a number of well defined security zones, network segmentation can reduce the impact of lateral movement, support regulatory compliance and improve cyber resilience in general. Regulations and standards, such as the NIS and NIS2 directive [1], DORA [2], the Cyber Resilience Act [3], and BSI IT Grundschutz [4] explicitly mention the need for structured network separation. However, network segmentation remains difficult to implement and maintain in practice [5], [6]. Large enterprise networks are often historically grown, heterogeneous and are if at all, only partly documented. Existing documentation may not accurately show the effective security zones that are enforced by network devices e.g. firewalls. Traditional approaches for understanding segmentation, manual documentation reviews, configuration inspections or administrator interviews require a lot of resources. Furthermore they are error prone and do not scale with the growing size of infrastructures. Network enforcement points continuously generate large volumes of traffic logs that contain information about real data traffic. These logs represent a rich, but so far, unexploited source of information about how network segmentation is enforced at runtime. Hence, the research question investigated in the ongoing PhD project is, how and to which extent real world network security zones and further structural network information can be reconstructed solely from log data, without relying on existing network documentation or configuration knowledge while producing consistent and scalable results across large IT landscapes. Existing work focuses on identifying statistical relationships in time series data [7], or on detecting temporal irregularities using machine learning approaches [8]. Both scenarios do not recover network segment structures. In addition, many methods rely on predefined network segments as input for further processing. This gap in literature motivates to introduce a log-based recovery approach of network segments. As a first step, a log-based, metadata-driven methodology that reconstructs security zones by iteratively analyzing the communication patterns observed at enforcement points is introduced. The approach focuses on network traffic analysis at open systems interconnect model (OSI) layer 3 and applies network modeling and measurement techniques. The applicability of the approach is demonstrated by evaluating it with log data of real companies' network architectures.

The remainder of this article is structured as follows. Section II summarizes background information and provides an overview of related work. Section III presents the proposed algorithm, while Section IV gives some evaluation results. Section V provides conclusions and an outlook to future work.

II. Background and Related Work

In this work a network segment is defined as a set of Internet Protocol addresses (IPs) that can communicate with each other without passing a network security device. Those segments are extracted out of predefined scopes e.g. RFC1918 blocks [9] that define the maximum possible, non-overlapping range of addresses for IT network communication participants. A network scope does not contain any information about the relation of the extracted network segments and their usage. The definition of network segments and their proper sizing within an IT landscape is vital for security reasons, as in case of network compromise, unrestricted and unintended communication within an enclave is possible. This

ISBN 978-3-903176-82-9

communication may lead to severe consequences e.g. data loss and business disruption. Existing research on network segmentation focuses primarily on design and optimization of large network architectures. This improves security as smaller and more isolated network zones significantly reduce the impact of compromises and lateral movement [6]. Redesign of existing networks into more secure architectures has been proposed via restructuring enclaves, refactoring topologies, or applying compartmentalization and layered protection [5], [10], [11]. Probabilistic and vulnerability aware models, such as continuous-time Markov chains, have been introduced to analyze how segmentation and patch cycles influence attack propagation [12]. Referential penetration testing approaches compare real environments against idealized segmentation policies to identify weaknesses [13]. While these methods provide insights how networks should be structured, they typically assume that the current network topology and segmentation are already known and documented. Yet, enterprise networks evolve over time and thus documentation about network segements mostly is incomplete and/or outdated and collecting this information by human labour is time-consuming, expensive and error-prone. Other approaches use graph reconstruction to identify time dependent relationships between networks [7] but do not reconstruct network enclaves itself. Statistical approaches reconstruct missing links and weights from aggregate information by the use of probabilistic models [14]. A structural knowledge of the network is mandatory, which emphasizes the need for a log-based reconstruction algorithm. Machine learning-based algorithms focus on anomaly detection of sequential data [8] instead of reconstructing security enclaves. Therefore, the PhD project attempts the reconstruction of effective segmentation in existing real-world networks based purely on observed traffic metadata, to enable later validation or optimization.

III. Methodology

The methodology focuses on Internet Protocol version (IPv)4 based enterprise networks and operates on OSI Layer 3 traffic metadata extracted from enforcement points such as firewalls. The key unit of analysis is a *Connection Pair* (CP) that is defined by a source IP and a destination IP address observed in a routed communication event.

If a communication event between two IPs is observed in enforcement point logs, the source IP and the destination IP cannot belong to the same network segment. Otherwise the traffic would not be routed and therefore not get logged. RFC1918 private address ranges [9] i.e. 10/8, 172.16/12, 192.168/16 are mostly used to design enterprise networks. This approach requires only a predefined abstract network scope that represents the maximum address spaces within real networks. The methodology does not assume any prior knowledge of subnetwork (SN) dividers within these scopes.

Certain scenarios are currently not considered, including overlapping IP ranges, the use of Network Address Translation (NAT) for East-West traffic and IPv6 environments. In most scenarios NAT is not used for East-West communication (internal communication), instead it is only used for internet-facing communication. As the algorithm benefits from analyzing internal communication it can be applied in most on-premise, hybrid and cloud environments. The objective is not to reconstruct physical or logical network designs exactly but to approximate security zones as enforced by routing and filtering behavior to complement and update existing documentation.

The proposed solution incrementally constructs an abstract network model as new CPs are processed. Each iteration improves the approximation of real world security zones by splitting larger address spaces into smaller SNs whenever required by observed traffic behavior. The methodology entails a set of statements, of which each represents a specific classification or modification step. This design ensures deterministic behavior and supports consistent results even when log entries are processed in different orders. The procedure of the algorithm is presented in Fig. 1.

A. Network Scope Classification

For each CP the source and destination IPs are evaluated against all predefined network scopes. The following two cases are distinguished.

1) Dual Scope Classification: Both CP IPs are within the same network scope. In this case, further refinement is required because the observed communication implies that the scope contains at least two distinct real networks.

2) Single Scope Classification: The CP IPs belong to different network scopes. Separation is already implied by the boundaries of the scope. A split is not necessary at the scope level. CP IPs that cannot be mapped to any scope do not contribute useful information and are excluded from further processing.

B. Super Network Classification and Network Splitting

If refinement within a network scope is required, super network classification derives the smallest SNs that ensure communicating IPs are placed in distinct network segments. This process incrementally refines the model by partitioning larger address spaces only where necessary. Two split operations are used.

1) Dual Split: If both IPs are part of the same smallest matching super network, that network is split into two smaller SNs. The split location is calculated using the binary difference between the IPs, to ensure minimal and deterministic segmentation.

2) Single Split: If only one IP is not covered by an existing SN of a network scope, while the other IP is already assigned, a new SN is created within the corresponding super network for the uncovered IP. The SN is selected with the largest possible non-overlapping address range within the available space.

In scenarios with incomplete SN coverage or gaps between existing SNs, splitting is applied selectively to uncovered areas. This may result in the creation of new SNs for specific IPs, followed by additional refinement steps if both IPs are

subsequently assigned to the same, newly created SN. These mechanisms ensure consistent address space coverage, prevent overlaps and support the iterative refinement of the network model.

C. Hit Markers and Model Compilation

A *Hit Marker* (HM) is an additive, that marks observed communication and its occurrence as subset of a matching network scope. HMs ensure that only network scopes with observed CPs are part of the final model. Each IP observed in a CP is associated with a HM, which is stored in the smallest matching network segment that includes the corresponding IP. When networks are split, HMs are migrated to the appropriate child networks if possible. Additional rules ensure that HMs assigned to a parent network are iteratively reevaluated and migrated to newly created SNs as they become available, preventing inconsistent assignments of HMs to parent networks. If a SN contains both, a HM and a more specific child SN, and the migration of the HM is not possible due to incomplete SN coverage (i.e. gaps), a new SN is created to accommodate the remaining HMs. This ensures that no parent network with a HM overlaps with child networks that also contain HMs.

The final model is compiled by collecting all SNs that contain at least one HM. The result is the list of found SNs where at least one HM is set. These SNs represent the actual security zones that are in use.

IV. Evaluation

An evaluation of the algorithm is conducted using synthetic data, as well as real world data of two organizations to demonstrate its applicability as well as the robustness of the algorithm. First, synthetic data is used to validate the functional correctness and permutation invariance of the model. Different orders of CPs keep constructing consistent security zones.

Second, the algorithm is applied to two distinct enterprise IT landscapes of which the actual segment structure had been manually compiled. The firewall logs were collected from two large company networks, which remain anonymous for confidentiality and security reasons. Both IT landscapes include multiple firewall systems and segmented internal address spaces. The network structures were reconstructed on expert knowledge and configuration data. CPs of firewall systems are used to construct effective security zones. Focus on model growth and approximation quality is set, considering the relationship between discovered IPs and reconstructed network segments. The reconstructed segments were compared to a manually derived ground of truth in both environments. The result shows that in both scenarios, a significant number of zones can be recovered early. Minor deviations occur mainly due to incomplete log coverage. Permission to use the data for this particular scientific purpose has been granted by both organizations. In both environments a non-linear, saturating growth is identified, illustrating the effectiveness of the algorithm in early stages of processing. The average network size converges as more data is processed, indicating increased model stability. A representation of organization A is shown in Fig. 2. While the organizations are different in size and structure and thus reveal differences in data saturation and segmentation granularity, the overall outlook is similar so that organization B is not shown.

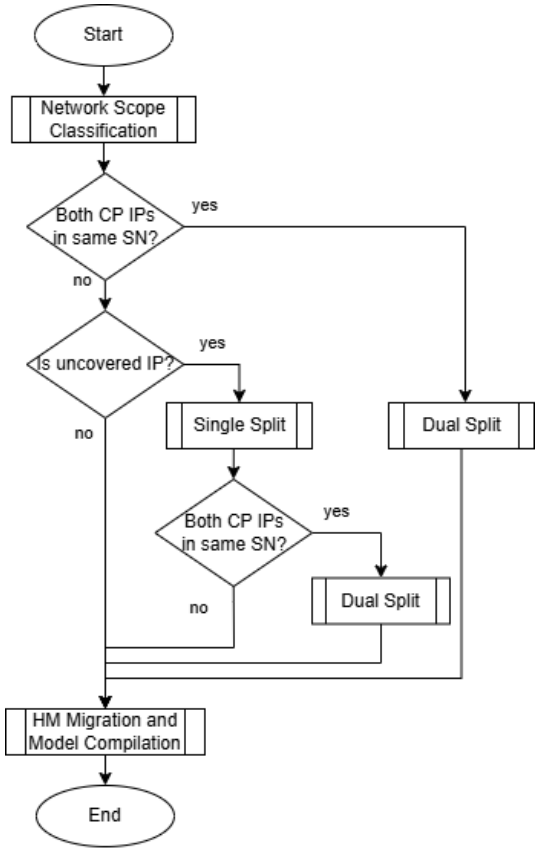

Fig. 1. Flowchart of the proposed approximation methodology. Super Network Classification and Network Splitting is visualized in detail.

During the process, reconstructed networks based on source IPs are more rapidly growing at the beginning, followed by an intersection point. The gradient of network growth between source and destination IPs shows a divergence at the end. This indicates multiple sources from different zones communicate to destinations within less divergent zones. This behavior changes at the end of the analysis.

The results demonstrate permutation invariance, robustness across different input orders, and practical applicability in heterogeneous IT environments.

V. Conclusion and Future Work

We presented an algorithm for network modeling and security architecture analysis. Evaluation demonstrated, that approximation of real-world network security zones could be derived solely from enforcement point log data, without relying on documentation or configuration access. The methodology is scalable, permutation invariant and applicable to large enterprise environments.

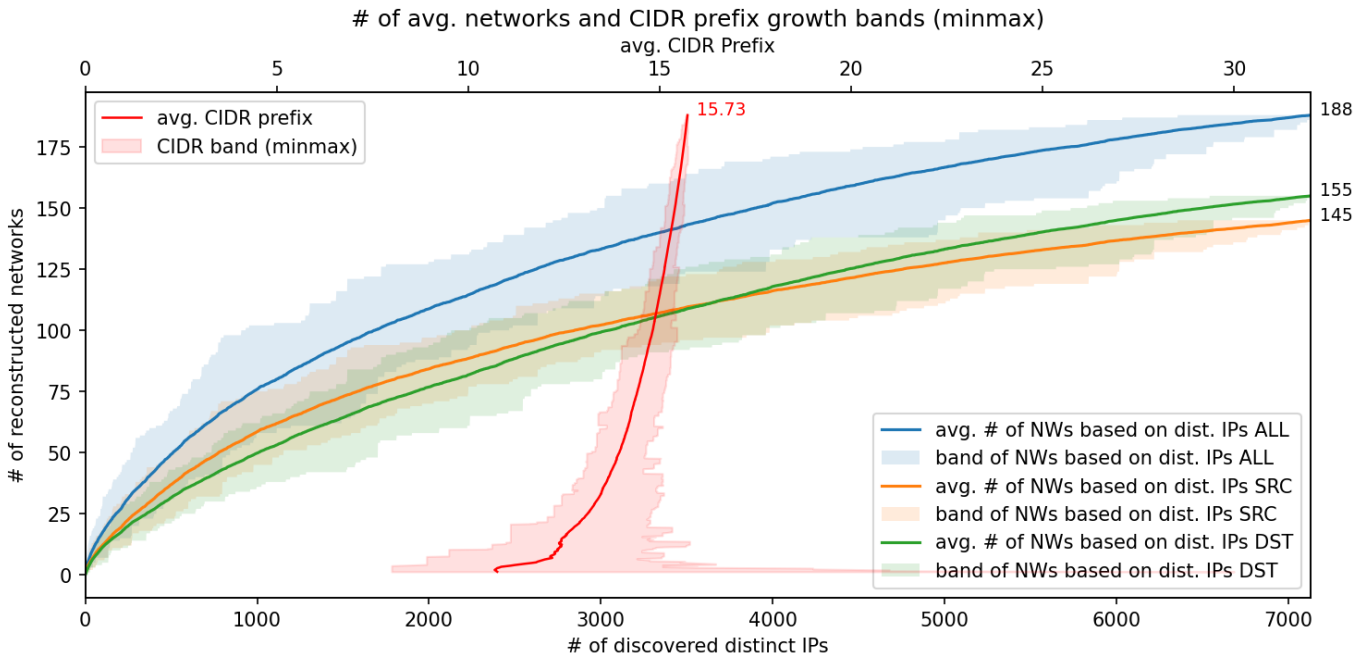

Fig. 2. Iteration plot of model reconstruction of the dataset from organization A with 100 permutations of CP order.

Future work will address limitations related to NAT, overlapping address spaces, and IPv6, using the algorithm with other log data as well as applying it in network traffic analysis for performance optimization and extended architectural coverage.

References

[1] European Union, "Directive 2022/2555," [accessed 2025-11-05]. [Online]. Available: https://eur-lex.europa.eu/eli/dir/2022/2555/oj/eng

[2] ——, "Regulation 2022/2554," [accessed 2025-11-05]. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2022/2554/oj/eng

[3] ——, "Regulation 2024/2847," [accessed 2025-11-05]. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2024/2847/oj/eng

[4] Bundesamt Sicherheit in der Informationstechnik, "BSI-Standard 200-2: IT-Grundschutz-Methodology," [accessed 2025-11-06]. [Online]. Available: https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Grundschutz/International/bsi-standard-2002_en_pdf.html?nn=128640

[5] N. Wagner, C. S. Sahin, J. Pena, and W. W. Streilein, "A nature-inspired decision system for secure cyber network architecture," in *2017 IEEE Symposium Series on Computational Intelligence (SSCI)*, 2017, pp. 1–8.

[6] E. Hemberg, J. R. Zipkin, R. W. Skowyra, N. Wagner, and U.-M. O'Reilly, "Adversarial co-evolution of attack and defense in a segmented computer network environment," in *Proc. Genetic and Evolutionary Comput. Conf. (GECCO '18)*. New York: ACM, 2018, pp. 1648–1655.

[7] M. Dimovska and D. Materassi, "An efficient network reconstruction method and applications," *IFAC-PapersOnLine*, vol. 54, no. 7, pp. 49–54, 2021, 19th IFAC Symposium on System Identification SYSID 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2405896321011071

[8] C. Zhang, G. Li, Y. Qi, H. Ye, L. Qing, M.-H. Yang, and Q. Huang, "Dynamic erasing network with adaptive temporal modeling for weakly supervised video anomaly detection," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 36, no. 9, pp. 16 706–16 720, 2025.

[9] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. Doering, and E. Gerich, "Rfc 1918: Address allocation for private internets." [Online]. Available: https://www.rfc-editor.org/rfc/rfc1918

[10] N. Mhaskar, M. Alabbad, and R. Khedri, "A formal approach to network segmentation," *Computers & Security*, vol. 103, p. 102162, 2021.

[11] M. Alabbad, N. Mhaskar, and R. Khedri, "Two formal design solutions for the generalization of network segmentation," *Journal of Network and Computer Applications*, vol. 222, p. 103763, 2024.

[12] N. Wagner, C. c. Şahin, J. Pena, J. Riordan, and S. Neumayer, "Capturing the security effects of network segmentation via a continuous-time Markov chain model," in *Proc. 50th Annual Simulation Symp.*, 2017.

[13] M. Alabbad, N. Mhaskar, and R. Khedri, "Hardening of network segmentation using automated referential penetration testing," *Journal of Network and Computer Applications*, vol. 224, p. 103851, 2024.

[14] G. Cimini, R. Mastrandrea, and T. Squartini, *Reconstructing Networks*. Cambridge University Press, Aug. 2021. [Online]. Available: http://dx.doi.org/10.1017/9781108771030