

Towards a Zero-Trust Micro-segmentation Network Security Strategy: An Evaluation Framework

Nardine Basta
Macquarie University
nardine.basta@mq.edu.au

Muhammad Ikram
Macquarie University
muhammad.ikram@mq.edu.au

Mohamed Ali Kaafar
Macquarie University
dali.kaafar@mq.edu.au

Andy Walker
ditno
Andy@ditno.com

Abstract—Micro-segmentation is an emerging security technique that separates physical networks into isolated logical micro-segments (workloads). By tying fine-grained security policies to individual workloads, it limits the attacker’s ability to move laterally through the network, even after infiltrating the perimeter defences. While micro-segmentation is proved to be effective for shrinking enterprise networks attack surface, its impact assessment is almost absent in the literature. This research is dedicated to developing an analytical framework to characterise and quantify the effectiveness of micro-segmentation on enhancing networks security. We rely on a twofold graph-feature based framework of the network connectivity and attack graphs to evaluate the network exposure and robustness, respectively. While the former assesses the network assets connectedness, reachability and centrality, the latter depicts the ability of the network to resist goal-oriented attackers. Tracking the variations of formulated metrics values post the deployment of micro-segmentation reveals exposure reduction and robustness improvement in the range of 60% – 90%.

I. INTRODUCTION

Micro-segmentation [7] is a core component of the zero-trust security concept. It creates secure zones across cloud and data centre environments to isolate the different application workloads and secure them independently. It further generates dynamic access control policies that limit network and application flows between workloads. Accordingly, it protects the network assets and provides control and visibility over the growing amount of east-west traffic across the organization which bypasses the traditional firewalls.

Autonomously modelling application behaviour and accounting for workloads is a major challenge towards achieving a zero-trust architecture. This can be attributed to the fact that enterprises data centre has evolved from on-premises infrastructure to a distributed facility with inter-connected cloud infrastructure where networks, applications and workloads are virtualized in multiple private and public clouds. Hence, in addition to the introduced complexity of the network architecture, enterprises have little confidence in their underlying network structure and connectivity.

While several micro-segmentation solutions are currently offered by industry (e.g., *ditno* [25], Cisco [9], and others), the general question of how effective and efficient these security controls are, still persists. In fact, little is known about how implementing these controls would compare to security risks within flat networks. This is not only essential to understand what level of protection is offered by the different

security controls, but is also important to justify resources and investment to augment existing controls.

In this paper, we leverage attack-graph generation and probabilistic reasoning framework for comprehensive security and effectiveness analysis of network micro-segmentation. The main contributions of our work are as follows:

- We develop a framework to assess and quantify the effectiveness of micro-segmentation in reducing the enterprise network assets risk of exposure to insider and outsider threats. We further analyse the robustness of the network by measuring its ability to resist goal-oriented attackers.
- To generate a reproducible and objective evaluation framework we base our metrics on graph feature analysis of the network connectivity and attack graphs.
- We rely on two enterprises network data to perform an empirical analysis of the effectiveness of micro-segmentation in light of the formulated framework. To the best of our knowledge, no previous work to quantify the impact of micro-segmentation, while relying on real-life network traffic, exists.
- Using data-sets from two real-world enterprise networks, we show that micro-segmentation successfully doubles the chain an attacker is forced to pursue to compromise a target network asset by automatically identifying and blocking the illegitimate network internal connections.
- While micro-segmentation is unable to reduce the network vulnerabilities, we show that modifying the system security configurations influences the likelihood of exploiting the vulnerabilities. We demonstrate that micro-segmentation decreases the network misconfigurations by 65% and the number of possibilities an intruder can exploit the network vulnerabilities by 99%.
- Additionally, we show that micro-segmentation contributes to enhancing the visibility of the network architecture. By identifying and classifying the network applications/services workloads, it enables highlighting the network misconfigurations and illegitimate connections.
- Our centrality metrics provide insight into the network weak links that should be prioritised for redemption.

II. BACKGROUND AND RELATED WORK

A. Background

Micro-segmentation is an implementation of a distributed virtual firewall that regulates access to network assets based

on security rules that have been determined on each workload (micro-segment). The firewalls examine the internal network traffic up to layer four (transport layer) of the Open Systems Interconnection (OSI) model and enforce access control through the generated micro-segmentation policies. Figure 1 depicts the connectivity structure of a sample enterprise network before and after micro-segmentation.

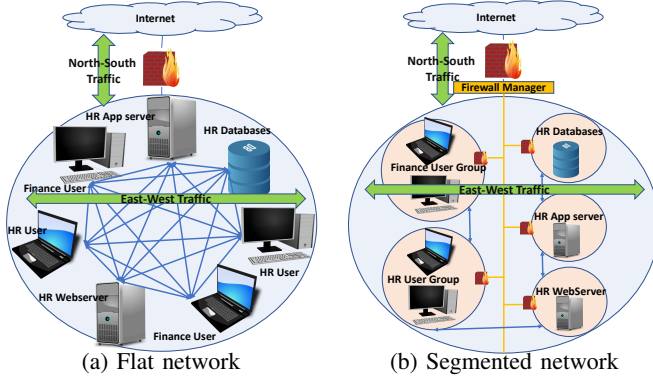


Fig. 1: An example of flat and micro-segmented network topologies of an enterprise network. In Figure 1a, the network assets are fully connected with no restrictions on the internal communication. In Figure 1b, the network is micro-segmented into different workloads (shown in circles), thus restricting the east-west traffic.

Unlike traditional firewalls that have no business logic or context into applications, the micro-segmentation firewalls are often implemented in software and deployed on each network workload. Machine-learning techniques are used to leverage the meta-data from micro-segments for firewall rules generation and generalisation across similar workloads. The network micro-segmentation firewalls are all linked to central management that pushes the policies to each firewall. It, therefore, enables granular policy enforcement throughout the enterprise network, not just at the perimeter.

There exist different possible architectures for dividing IT services into multiple tiers. Services are most commonly divided into three tiers; web servers, application servers and databases. Nevertheless, some approaches combine the web server and application server tiers. Workloads within a single IT service are allowed to communicate to each other through very restrictive rules to minimise the propensity of lateral movements. As a result, micro-segmentation offers better visibility and accountability of network resources and communication among micro-segments.

B. Related Work

Efforts have been deployed in the domain of measuring and quantifying network security in general and the effectiveness of firewalls in particular. In fact, quantifying security is one of the well-recognized open problems [18], [22], [23] where most of the approaches quantify the overall network security performance [4], [14], [15], [23], [29]. Some approaches are specially dedicated to analysing the deployed firewalls [5] and quantifying their effectiveness [6], [8], [28].

Although micro-segmentation has gained momentum in reducing the network attack surface, improving breach containment and strengthening regulatory compliance, the assessment and quantification of its performance are almost absent in the academic and industrial literature. Only one approach is found to quantify the overall efficacy of micro-segmentation by measuring the time it takes the attacker to traverse the network and compromise its assets [17]. Nevertheless, the utilised metric is based on the attacker’s skills rather than the strength of the micro-segmentation security controls. Additionally, the tests were conducted in a simulated network environment rather than a real network.

In this paper, we fill the gap by providing more comprehensive and reproducible metrics to accurately represent the impact of micro-segmentation on enhancing the security of enterprise networks.

III. MICRO-SEGMENTATION IMPACT QUANTIFICATION FRAMEWORK

Next, we outline our framework which assesses the impact of micro-segmentation on enterprises network security through quantitative metrics based on graph analysis. In particular, we characterise the impact of micro-segmentation on the network *exposure* and *robustness*.

A. Network Exposure

Keeping pace with the dynamic nature of the industry, enterprise networks are evolving fast. In fact, network exposure, due to expanding the network connectivity, is the new IT risk many businesses are ignoring at their peril. Hence, comes the importance of micro-segmentation to block the illegitimate, insecure or unneeded connections and control the east-west traffic within the enterprise network.

To formally define network exposure, let $C(A, V, w)$ be the network connectivity directed graph where V is the set of graph vertices representing the network assets. A is the set of graph directed edges indicating that the connected vertices are allowed to communicate, where $A \subseteq \{(x, y) \mid (x, y) \in V^2 \wedge x \neq y\}$. $w : A \rightarrow \mathbb{R}$ is a weight function representing the number of services that the network assets are allowed to use for communicating in the direction of the connection. We classify the network exposure metrics into three categories: *connectedness*, *reachability*, and *centrality*.

Connectedness: A direct indicator of the internal network exposure is the amount of connections allowed to and from the different enterprise network assets. In fact, the higher the connectedness between the network assets, the more the possibilities/ways that are presented to attackers in order to achieve their malicious target.

We formulate the first metric to quantify the enterprise network exposure in this context, namely the Enterprise Network Internal Connectivity Exposure (ENICE). We compute the ENICE metric following Equation 1 where $w(a)$ is the weight of edge $a \in A$.

$$ENICE = \sum_{a \in A} w(a) \quad (1)$$

The clustering coefficient of a connectivity graph measures how interconnected a vertex's neighbours are to one another. The higher the clustering coefficient, the more exposed the network assets are to lateral attacker movement. The global clustering coefficient is designed to give an overall indication of the clustering in the network based on triplets of nodes. A triplet is three nodes that are connected by either two (open triplet) or three (closed triplet) ties. Accordingly, the second exposure metric, namely the Global Clustering coefficient (GC), is calculated as follows:

$$GC_C = \frac{\text{number of closed triplets}}{\text{total number of triplets}} \quad (2)$$

Reachability: Given the network connectivity graph, finding out the vertices that are reachable from another vertex v is an indicator of the number of possibly compromised assets in case v is illegitimately accessed. A direct measure of reachability is the shortest path analysis of the network connectivity graph. The shortest path represents the minimum number of hops an attacker needs to exploit until reaching a target asset. The longer the shortest path, the more effort an attacker needs to deploy to compromise a target network asset. The Mean of shortest Path Length (MPL) represents the average number of hosts, in the best case, an attacker needs to compromise in order to reach their target. It is a strong indicator of how likely are the network assets reachable from a compromised source. Let us denote by LSP_C the list of shortest paths in the connectivity graph C and by p a path in the connectivity graph where $|p|$ is the number of vertices in the path p . MPL_C is calculated following Equation 3.

$$MPL_C = \frac{1}{|LSP_C|} \sum_{p \in LSP_C} |p| \quad (3)$$

The diameter of a graph is the maximum eccentricity of any vertex in the graph. It is equal to the maximum shortest path length of the connectivity graph. It indicates the longest chain an attacker is forced to pursue to compromise a target network asset from a compromised source. The latter assumes that attackers usually aim at taking the shortest path to reach their target asset. Let $s(u, v)$ be the shortest path distance from vertex u to vertex v . The Connectivity graph Diameter (CD), the second reachability metric, can be defined by:

$$CD_C = \max_{v, u \in V} \{s(u, v)\} \quad (4)$$

The transitive closure of a connectivity graph C is a graph $C^T = (A^T, V)$ such that for all $i, j \in V$ there is a link (i, j) if and only if there exists a path from i to j in C . Therefore, it provides a more detailed perspective of how far an attacker can go after compromising a network asset. A direct indicator of the latter is the number of transitive paths in the network represented by the transitive closure graph edges. Standing on this ground, we proceed by formulating the third reachability metric, namely the Transitive Internal Network Reachability (TINR), following Equation 5 where A^T is the set of edges of

the transitive closure graph C^T corresponding to the network connectivity graph C .

$$TINR_C = |A^T| \quad (5)$$

Centrality: Nodal centrality quantifies how important a node is within a network. Hence, it implies the impact of compromising the node on the overall network security. The out-degree centrality metric of the network connectivity graph vertex v outlines the number of possibly compromised assets by the attacker in case v is compromised. Accordingly, we develop a metric namely AVerage Out-Degree (AVOD) to account for the typical number of possibly compromised assets after a successful attack. Let $OD(v)$ be the out-degree of a connectivity graph $C(A, V)$ node $v \in V$. The average out-degree is given by:

$$AVOD_C = \frac{1}{|V|} \sum_{v \in V} OD(v) \quad (6)$$

The closeness centrality is tightly related to the notion of distance between nodes. It highlights nodes that may reach any other nodes within a few hops and nodes that may be very distant. The closeness centrality of a network asset v indicates how fast an attacker, after compromising v , can access all other nodes in the network. The Closeness centrality (CL) of node v is computed following Equation 7 where $d(v, u)$ is the distance (number of vertices) between v and u :

$$CL(v) = \frac{1}{\sum_{u \in V} d(v, u)} \quad (7)$$

The average closeness of the network connectivity graph nodes, the second centrality metric, is an indicator of how fast the whole network can be compromised after a breach. It is calculated as follows:

$$AC_C = \frac{1}{|V|} \sum_{v \in V} CL(v) \quad (8)$$

B. Network Robustness

Network robustness is the ability of a network to resist goal-oriented attackers. To measure the network robustness, we leverage the attack graph structure in conjunction with a component metric such as the Common Vulnerability Scoring System (CVSS) [16], [24]. An attack graph is a succinct representation of all paths through a system that ends in a state where an intruder has successfully achieved his goal [13]. We classify the attack graph-based network robustness security metrics into three categories; non-path-based, path-based and CVSS-based security metrics. We generate the attack graph using the MulVAL tool [19]–[21] while relying on the Nessus [3] vulnerabilities scanner output and the enterprise network perimeter and micro-segmentation firewall rules.

To formally define network robustness, let $G(E, N)$ be the enterprise network attack graph consisting of a set of nodes N of three types: attack step nodes, privilege nodes and configuration nodes [10]. Let R_G be the root nodes of the attack graph G representing network configurations that

contribute to attack possibilities. Let L_G be the set of privilege nodes denoting the compromised assets. The set of paths P_G of the attack graph G comprises all directed attack paths starting at the root configuration nodes R_G and ending at the privilege nodes L_G .

Path-Based metrics: The shortest path metric indicates the minimum number of attack steps an attacker should perform to compromise an asset in the network. Since a chain is only as strong as its weakest link, it is a significant indicator of the network robustness. Indeed, attackers would need a global view of the system vulnerabilities to deliberately exploit the shortest path. However, the probability of the shortest path to be exploited is directly proportional to the number of shortest paths in the network [12]. Hence, we define the Number of Shortest Paths metric (NSP) to identify the count of shortest attack paths between every root node $r \in R_G$ and privilege node $l \in L_G$ in the attack graph G . Let us denote by $LSAP_G$ the list of shortest attack paths in the attack graph G . The *NSP* metric is given by:

$$NSP_G = |LSAP_G| \quad (9)$$

The Minimum Shortest Path Length metric (MSPL) denotes the absolute minimum number of attack steps an attacker needs to perform to compromise a target network asset. Let p be a path in the attack graph where $|p|$ is the number of nodes in p , accordingly:

$$MSPL_G = \min_{\forall p \in LSAP_G} \{|p|\} \quad (10)$$

The count of paths with length equal to the minimum shortest path length outlines the number of network weakest links. Hence, in conjunction with the MSPL metric, it is a direct indicator of the likelihood of a successful attack. Consequently, we define the Count of Minimum Path Length (CMPL) metric. Let LMP_G be the list of paths with length equals to $MSPL_G$ in the attack graph G . Accordingly, *CMPL* is given by:

$$CMPL_G = |LMP_G| \quad (11)$$

Non-Path-Based metrics: The count of the attack graph configuration nodes is an intuitive yet significant non-path-based measure of the network robustness. The configuration nodes depict facts about the current network configuration that contributes to one or more attack possibilities. Hence, they stand as the attacker's entry point to the network. We proceed by formulating the first non-path metric namely the Count of Mis-Configurations (CMC) following Equation 12 where R_G is the set of root nodes of the attack graph G :

$$CMC_G = |R_G| \quad (12)$$

The out-degree metric (OD) of the attack graph privilege nodes L_G is an indicator of the number of possible attacks succeeding a compromised network asset. Hence, we define the average out-degree AOD_G and the maximum out-degree MOD_G metrics. While the earlier is an indicator of the typical number of attacks after a compromised asset the latter

represents the worst-case scenario, indicating the maximum number of possible attacks.

Let $OD(n)$ be the out-degree of an attack graph privilege node n . The Average Out-Degree (AOD) and the Maximum Out-Degree (MOD) metrics of an attack graph G are given by Equations 13 and 14, respectively.

$$AOD_G = \frac{1}{|L_G|} \sum_{\forall n \in L_G} OD(n) \quad (13)$$

$$MOD_G = \max_{\forall n \in L_G} \{OD(n)\} \quad (14)$$

The betweenness centrality of an attack privilege node is the extent to which the vertex plays a bridging role in a network. In other words, it measures the extent that the privilege node falls on the shortest path to other privilege nodes. Accordingly, the higher the betweenness of a compromised privilege node, the more widespread across the network the subsequent privileges the attacker can acquire. The betweenness metric further provides insight into the network weak links that should be prioritised for redemption. The betweenness centrality of a privilege node n is calculated following Equation 15 where NSP_{rl} is the number of shortest paths from root $r \in R_G$ to privilege $l \in L_G$, $NSP_{rl}(n)$ is the number of shortest paths passing through $n \in L_G$ and $r \neq l \neq n$.

$$BN(n) = \sum_{\forall r \in R_G \wedge \forall l \in L_G} \frac{NSP_{rl}(n)}{NSP_{rl}} \quad (15)$$

The Average Betweenness (AB) of the attack graph privilege nodes depicts the typical contribution of an illegitimately acquired network privilege in other potential attacks across the network compromising privileges that the attacker could not have reached otherwise. It is calculated following Equation 16 where L_G is the set of privilege nodes of the attack graph G .

$$AB_G = \frac{1}{|L_G|} \sum_{\forall l \in L_G} BN(l) \quad (16)$$

Common Vulnerability Scoring System (CVSS) metrics: Assigning complexity values of the network vulnerabilities communicates their characteristics and severity. This approach reflects variations in the difficulty of exploiting the different vulnerabilities [27]. A standard, such as the CVSS [24], may be used to provide guidance in scoring vulnerabilities.

While there is currently no standard way of aggregating vulnerability metrics, a critical issue in measuring network security is to combine measures of individual vulnerabilities, and configurations into a global measure [11]. We rely on the work presented in [11], [26], [30] to analyse the enterprise network attack graph for the purpose of calculating cumulative metrics and aggregating the vulnerabilities score.

The cumulative score of a given privilege node indicates the likelihood that the corresponding resource is compromised during an attack, or equivalently, among all attackers attacking the network over a given time period, the average fraction of attackers who can successfully compromise the resource,

taking into consideration the effects of all possible interplays between vulnerabilities. We refer the reader to [26] for further details on the underlying mathematical modelling and computation of risk.

IV. ANALYSIS AND RESULTS

In this section, we assess the effectiveness of the proposed framework in evaluating the impact of micro-segmentation on enterprise network security. To construct the network connectivity graph, representing the basis of the exposure analysis, we identify the network hosts forming the graph nodes. While the flat network connectivity graph is a complete graph, we rely on the micro-segmentation firewall rules to identify the edges of the segmented network connectivity graph. The firewall rules are described in terms of the source host, destination host, service protocol and service destination port.

To assess the network robustness, the attack graph is generated using the MulVAL tool. It is given as input the Nessus XML output and the network firewall rules. While the earlier depicts the hosts' vulnerabilities and configurations including the running software and services, the latter models the assets access control policies regulating their communication. The firewall rules include the micro-segmentation rules, if any, in addition to the enterprise network firewall rules regulating north-south traffic. In the case of a flat network, the micro-segmentation rules are replaced by one generic rule allowing all internal network traffic.

We consider the network data of two enterprises; a university and a life-care organisation. Table I provides statistics describing the two networks. The connections depict the count of asset pairs that are allowed to communicate. While we were able to assess the exposure of both enterprises, due to the unavailability of the Nessus data for the life-care organisation (Enterprise B), the robustness analysis was only limited to the university enterprise network (Enterprise A).

The metrics are calculated prior to and post the deployment of micro-segmentation. The results are then compared, assessed and used to draw conclusions which lead to useful insights that follow in this section.


A. Network Exposure Analysis

Connectedness: After micro-segmentation deployment, the ENICE values of Enterprises A and B networks are decremented by 99.92% and 99.78% respectively as shown in Figure 2a. Since the flat network topology, by definition, allows all internal network communications, we assumed that devices can listen on any port. Hence, for the flat network topology, we considered that all connections have an equivalent weight of 65,535. It can be argued that hosts are not necessarily listening on every possible port. However, in the flat network topology, nothing can stop a network asset from listening on a particular port. In fact, not many organizations, prior to deploying micro-segmentation, are fully aware of all the ports their assets are listening on. Therefore, micro-segmentation not only regulates the hosts that are allowed to directly communicate, it further

restricts the services used for communication between the different network assets.

Restricting the network connectedness further minimises the impact of the zero-day attacks by limiting the ports the micro-segments are listening on. Let us consider the zero-day threat caused by the Windows vulnerability SMBGhost [1] affecting the Microsoft Server Message Block (SMB) of the network file sharing protocol. Successful exploitation of this vulnerability enables the attacker to execute arbitrary code on the SMB server or client. The attacker could then possibly install programs, view, edit or remove data. Deploying micro-segmentation, the control rules deny the file-sharing capability on all the enterprise database servers while only allowing connectivity on the SQL port. Despite the presence of a zero-day vulnerability that could be exploited, no attacker is able to connect to the database server on the port that is vulnerable.

Similar behaviour is depicted by the global clustering coefficient. As shown in Figure 2b, for the flat networks, the GC metric values are close to one indicating a high percentage of closed triads in the graphs. The latter implies that the graph nodes are involved in as many transitive relations as possible. On the other hand, the micro-segmented network GC values are reduced by 80% and 70% for Enterprise A and Enterprise B, respectively. It can be remarked that Enterprise A clustering coefficient improvement percentage is higher than B. This can be attributed to the fact that B has leveraged infrastructure across applications thus allowing more transitive connections. In other words, Enterprise B has fewer dedicated unique servers per application. For example, it has three applications sharing the same database server.



(a) Network Internal Exposure (b) Global Clustering Coef.

Fig. 2: Connectedness metrics analysis.

Reachability: We begin by analysing the distribution of the shortest paths of the network connectivity graphs. Figure 3 presents the paths length distribution before and after micro-segmentation. The x-axis denotes the relation identifier of the possibly connected pair of hosts and the y-axis represents the shortest path length connecting the two nodes, if any. The shortest path length of the flat network of both enterprises has a constant value of one. Hence, an attacker can reach any other network asset in one step. This represents the situation of an employee laptop directly communicating with an enterprise database. In the event the laptop is infected with malware or CryptoLocker, it could directly infect the connected databases.

TABLE I: Summary of the dataset consisting of two networks' configurations

Enterprise	Business	# of Hosts	# Connections (Flat Network)	# of Connections (Segmented Network)
Enterprise A	University	300	90,000	4,045
Enterprise B	Life-care organization	238	56,644	3,007

Contrastingly, the majority of paths length in the segmented network of both enterprises falls in the range of $[2, 3]$ with an average value of 2.14 for Enterprise A and 2.17 for Enterprise B. It should be noted that the shortest path length values reflect the number of IT services tiers the organization have. For the two hereby considered enterprises, after micro-segmentation, services have either three tiers of workloads (user \rightarrow web server, web server \rightarrow application server, application server \rightarrow database) or two tiers (when the web server and application server tiers are combined). Consequently, the typical effort an attacker needs to deploy in order to reach a target asset is doubled as a result of micro-segmentation deployment.


The paths with length going to infinity characterising the segmented network of both enterprises designate that there is no path between the two possibly connected assets. The percentage of infinity paths of the segmented Enterprises A and B networks are 5.6% and 27.3% respectively which further limits the assets reachability by restricting the attacker's lateral movement within the network. It is worth mentioning that Enterprise A has fewer infinity paths because it has more management services characterised by their high connectivity.

Similar behaviour is observed with the worst-case reachability value depicted by the Connectivity graph Diameter metric (CD). Since the connectivity graph of the flat network is complete, the attacker can reach any target asset in exactly one step. Contrastingly, in the case of the micro-segmented network, the maximum effort deployed to attain the intended asset is tripled for both enterprises. This can be explained by the fact that, for these two enterprises, IT services are divided into a maximum of three tiers. Therefore, to compromise the database tier, the attacker needs to first gain access to the web server and the application server tiers.

Whilst the impact of micro-segmentation on the mean shortest path length and diameter of both enterprises is almost identical, the TNR metric exhibits a different behaviour as highlighted by Figure 4. For Enterprise A, the transitive reachability is decreased by 5.6% only whereas for Enterprise B, it is decreased by 27.3%. The latter values are found to be identical to the percentages of disconnected nodes as a result of deploying micro-segmentation. Hence, it can be safely deduced that the risk of compromising the network assets is a factor of the number of allowed internal network connections.

Centrality: The flat topology, by definition, allows all internal network traffic resulting in a complete connectivity graph. If any node is compromised the whole network is at risk as depicted by the linear out-degree distribution of Figure 5. In contrast, micro-segmentation significantly reduces the average out-degree centrality by 95.5% and 94.7% for Enterprise A and Enterprise B, respectively, as shown in Figure 7a.

Despite the undeniable improvement of micro-segmentation on the overall network out-degree centrality, Figure 5 exhibits



(a) University Network

(b) Life-care Network

Fig. 3: Connectivity graph shortest path length distributions.

unexpected behaviour of the values of the individual nodes. It can be noticed that some nodes have a high out-degree and can be a source of risk if compromised. After further analysis, they are found to belong to common management services (e.g. active directory, backup, ntp, dns, etc.) characterised by high connectivity. It should be noted that the hereby presented framework highlights, in different contexts, the network weak links that need to be given priority for investigation and mitigation.

The closeness centrality distribution values presented in Figure 6 fall in the range of $[0, 1]$. A value close to zero indicates that a given node is distant from other nodes in the network. It further signifies that numerous links need to be traversed to get to other nodes in the network.

It can be remarked that the flat closeness value for both enterprises is equal to one. This can be attributed to the fact that all flat network nodes are exactly one hop away distant from each other. Therefore, after acquiring an illegitimate privilege, an attacker can access any other asset by traversing only one link. On the other hand, after micro-segmentation, the overall closeness is reduced by almost half for both enterprises. The high spikes characterising the individual nodes' closeness distribution of Figure 6 belong to common management services.





Fig. 4: Transitive closure edges count.

B. Network Robustness Analysis

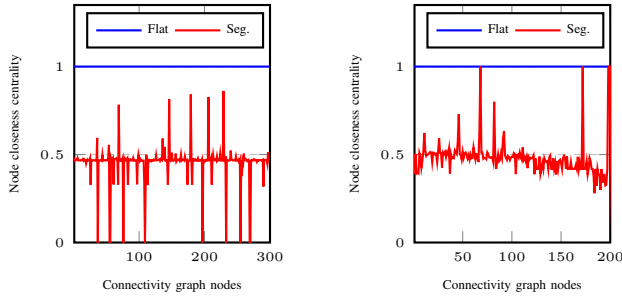
Path-based: Reducing the number of attack paths is a direct indicator of improving the network robustness reflecting the ability of networks to resist failures or attacks. Assessing the NSP metric representing the count of shortest attack paths, *before* and *after* micro-segmenting the network of Enterprise



(a) University Network

(b) Life-care Network

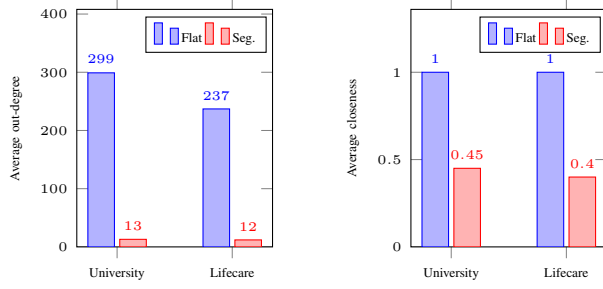
Fig. 5: Connectivity graph out-degree distribution.



(a) University Network

(b) Life-care Network

Fig. 6: Closeness centrality distribution.



(a) Average Out-Degree

(b) Average Closeness

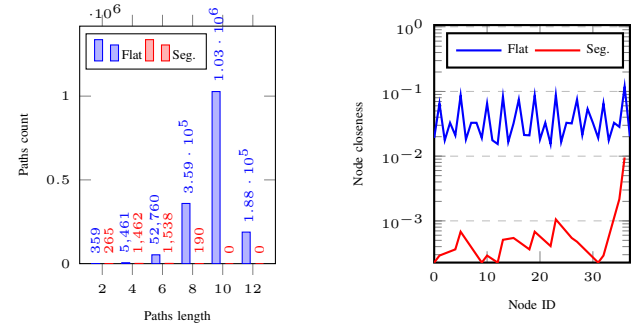
Fig. 7: Connectivity graph Centrality analysis.

A, reveals the efficacy of micro-segmentation to significantly reduce the total number of attack paths by 99%.

TABLE II: Robustness metrics values of Enterprise A.

Metric	Flat	Segmented	Improvement
Mis-configurations count	635	221	65.2%
Shortest paths count	1,633,440	3,455	99.7%
Average shortest path length	5	10	50%
Min shortest path length	2	2	0%
Min shortest path count	359	265	26.2%
Average out-degree	28	2	92.9%
Maximum out-degree	35	11	68.6%
Average betweenness	0.04	0.0005	98.8%

We proceed by investigating the distribution of attack paths lengths of Figure 8a. While the large number of paths with lengths greater than eight in the flat network might give the impression that the flat network is more resilient to attacks as



(a) Path Length

(b) Betweenness

Fig. 8: Attack graph paths length and betweenness distributions.

it requires more steps to be performed by the attacker, the total number of attack paths clears the doubt. Unlike the segmented topology, the unrestricted nature of the flat network exposes it to a multitude of direct and indirect potential attack vectors.

The MSPL metric output, representing the weakest link of the network, is the same for both the flat and micro-segmented networks. However, the count of paths with minimum length, indicating the probability that an attacker exploits the network's weakest link, is reduced by 26% after micro-segmentation deployment.

Non-path-based: Network misconfigurations raise the probability of success in exploiting the assets' vulnerabilities hence increasing the impact of the latter. While micro-segmentation is unable to reduce the vulnerabilities, we claim that modifying the system security configurations influences the likelihood of exploiting the vulnerabilities. We proceed by analysing the count of misconfigurations (CMC) metric. Although the number of network vulnerabilities has not changed as demonstrated by the Nessus scan, we notice that the number of the attack graph root nodes in micro-segmented network is reduced by 65%, as shown in Table II. Accordingly, we demonstrate that micro-segmentation significantly reduces the attacker accessibility to system vulnerabilities.

The AOD and MOD metrics results, indicating the typical and worst-case number of attacks succeeding a compromised network privilege, reveal the significant impact of micro-segmentation on improving the network robustness. It restricts the attacker's lateral movement and exploration, reducing the average and maximum possible attacks after a network breach by 93% and 69%, respectively, as indicated by Table II.

Finally, we analyse the betweenness centrality distribution of the network attack graphs before and after micro-segmentation. Compromised privilege assets with high betweenness centrality have a significant influence on the network robustness by virtue of their control over the paths leading to other privilege nodes, increasing the attacker ability to extend their attained privileges. In fact, micro-segmentation, not only reduced the average betweenness by more than 98%, it further changed the network topology resulting in a more linear distribution of the betweenness metric. Hence, unlike the jigsaw distribution provided by the flat topology which

TABLE III: Attack-graph nodes involving host V . The first row denotes the attacker acquired privilege on V . The second row presents the different vulnerabilities of V where each is identified by the NVD CVE unique identifier [24]. The third row enumerates V misconfigurations where the insecure connections are defined as hacl(Host1, Host2, Protocol/Service, Port).

	Flat Network Attack-graph Nodes	Micro-segmented Network Attack-graph Nodes
Privilege	execCode(V,user) 0.8	execCode(V,user) 0.64
Vulnerabilities	vulExists (V, CVE-2014-3802, debug interface access software development kit, remote-Exploit, privEscalation), vulExists (V, CVE-2017-5715, cortex-a, localExploit, privEscalation), vulExists (V, CVE-2017-5753, cortex-a, local-Exploit, privEscalation), vulExists (V, CVE-2017-5754, cortex-a, localExploit, privEscalation)	vulExists (V, CVE-2014-3802, debug interface access software development kit, remote-Exploit, privEscalation), vulExists (V, CVE-2017-5715, cortex-a, localExploit, privEscalation), vulExists(V, CVE-2017-5753, cortex-a, localExploit, privEscalation), vulExists (V, CVE-2017-5754, cortex-a, localExploit, privEscalation)
Configurations	hasAccount(V victim, V, user), networkServiceInfo(V, debug interface access software development kit, Windows, 2, user), hacl (internet, V, Windows, 2), hacl (V, AL, Windows Microsoft Bulletins, 2), hacl (V, AG, Windows Microsoft Bulletins, 2), hacl (V, AU, Windows Microsoft Bulletins, 2), hacl (V, AU, Windows, 2), hacl (V, BA, Windows Microsoft Bulletins, 2), hacl (V, AO, Windows, 2), hacl (V, HO, Windows Microsoft Bulletins, 2), hacl (V, KI, Windows Microsoft Bulletins, 2), hacl (V, KI, Windows, 2), hacl (V, MA, Windows Microsoft Bulletins, 2), hacl (v, MI, Windows Microsoft Bulletins, 2), hacl (V, MI, Windows, 2), hacl (V, NA, Windows, 2), hacl (V, NE, Misc, 3), hacl (V, NE, Windows, 2), hacl (V, OU, Windows, 2), hacl (V, OV, Windows Microsoft Bulletins, 2), hacl (V, PA, Database, 3), hacl (V, PA, Windows Microsoft Bulletins, 2), hacl (V, TE, Windows Microsoft Bulletins, 2), hacl (V, V, Windows, 2), hacl (V, WI, Windows Microsoft Bulletins, 2), hacl (V, WI, Windows Microsoft Bulletins, 3), hacl (TE, V, Windows, 2), hacl (OV, V, Windows, 2), hacl (NA, V, Windows, 2), hacl (MI, V, Windows, 2), hacl (MA, V, Windows, 2), hacl (KI, V, Windows, 2), hacl (HO, V, Windows, 2), hacl (CO, V, Windows, 2), hacl (BA, V, Windows, 2), hacl (AU, V, Windows, 2), hacl (AG, V, Windows, 2), hacl (AL, V, Windows, 2), hacl (WI, V, Windows, 2)	hasAccount (V victim, V, user), networkServiceInfo (V, debug interface access software development kit, Windows, 2, user), hacl (internet, V, Windows, 2)

is characterised by multiple peaks, it evenly distributed the reduced impact of compromising any asset on the network robustness as demonstrated by Figure 8b.

Common Vulnerability Scoring System (CVSS): Next, we analyse the quantitative security of a networked system through the cumulative probability that a network asset is compromised by an attacker. As previously discussed, the probability that an attacker succeeds in obtaining an illegitimate privilege is a function of the severity of a system vulnerability and the accessibility of the vulnerability to the attacker. While the first factor remains intact, after micro-segmentation the second factor is significantly altered as a result of modifying the network topology and limiting the accessibility to the network assets.

We proceed by calculating the cumulative risk values of the network privilege nodes before and after micro-segmentation following the model in [26]. The distribution of the calculated risk values reveals an average improvement of 20% after micro-segmentation deployment. While this value represents a significant improvement of the overall enterprise network security risk, it is remarkably lower than the calculated improvements achieved by assessing the network exposure and robustness. This can be attributed to the fact that a major parameter contributing to the risk calculated values, namely the vulnerabilities severity, identified by the Common Vulnerability Scoring System [2], remains unchanged.

To further analyse the latter finding, let us consider the attack graph privilege node **execCode(V, User)**. It implies that the attacker has gained user privilege to execute code on the victim host V . As shown in Table III, the cumulative risk

value of acquiring this privilege in the flat network is found to be 0.8 while in the micro-segmented network it is reduced to 0.64. Further analysis of the attack graphs confirms our interpretation of the aggregated risk values. Both attack graphs have the same vulnerabilities of host V (four) as represented by the Vulnerabilities row of Table III. On the other hand, the flat network has 38 misconfiguration nodes involving V while the segmented network comprises only three.

V. CONCLUSION

In this paper, we proposed a suite of metrics to analyse the impact of micro-segmentation on improving network security. We leveraged graph features to study the network exposure and robustness against attacks. Hence, the presented work is the *first* to formulate *objective* graph-features-based metrics for quantifying the effectiveness of micro-segmentation. We rely on real enterprise network data to perform an empirical analysis of the developed suite of metrics.

The analysis of the formulated metrics before and after the deployment of micro-segmentation proves that the latter is one of the most effective strategies to protect against cyber threats. Comparing and assessing the exposure and robustness metrics reveals an improvement in the range of 60% – 90%.

REFERENCES

- [1] “NATIONAL VULNERABILITY DATABASE (NVD): CVE-2020-0796,” <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2020-0796>.
- [2] “NATIONAL VULNERABILITY DATABASE (NVD): CVSS Vulnerability Metrics,” <https://nvd.nist.gov/vuln-metrics/cvss>.
- [3] “Tenable network security: The nessus security scanner,” <http://www.nessus.org>.

- [4] M. S. Ahmed, E. Al-Shaer, and L. Khan, "A novel quantitative approach for measuring network security," in *IEEE INFOCOM*, 2008.
- [5] E. Al-Shaer, H. Hamed, R. Boutaba, and M. Hasan, "Conflict classification and analysis of distributed firewall policies," *JSAC*, 2005.
- [6] M. N. Alsaleh, S. Al-Haj, and E. Al-Shaer, "Objective metrics for firewall security: A holistic view," in *CNS*, 2013, pp. 470–477.
- [7] A. C. S. Centre, "Strategies to mitigate cyber security incidents – mitigation details," <https://www.cyber.gov.au/acsc/view-all-content/publications/strategies-mitigate-cyber-security-incidents>.
- [8] H. Chen, J.-H. Cho, and S. Xu, "Quantifying the security effectiveness of firewalls and dmzs," in *HoTSoS*, 2018.
- [9] Cisco, "Cisco secure workload (tetration)," https://www.cisco.com/c/en_au/products/security/tetration/index.html.
- [10] J. Homer, S. Zhang, X. Ou, D. Schmidt, Y. Du, S. Rajagopalan, and A. Singhal, "Aggregating vulnerability metrics in enterprise networks using attack graphs," *Journal of Computer Security*, vol. 21, pp. 561–597, 07 2013.
- [11] —, "Aggregating vulnerability metrics in enterprise networks using attack graphs," *Journal of Computer Security*, 2013.
- [12] N. Idika and B. Bhargava, "Extending attack graph-based security metrics and aggregating their application," *IEEE TDSC*, 2012.
- [13] S. Jha, J. Wing, and O. Sheyner, "Two formal analysis of attack graphs," in *IEEE CSFW*, 2002.
- [14] X. Li, P. Parker, and S. Xu, "A stochastic model for quantitative security analyses of networked systems," *IEEE TDSC*, 2011.
- [15] —, "A stochastic model for quantitative security analyses of networked systems," *IEEE TDSC*, vol. 8, no. 1, pp. 28–43, 2011.
- [16] P. Mell, K. Scarfone, and S. Romanosky, "Nist interagency report 7435, the common vulnerability scoring system (cvss) and its applicability to federal agency systems," 08 2007.
- [17] R. NANDAKUMARA, "Efficacy of micro-segmentation assessment report," *Bishop Fox*, June 2020.
- [18] D. M. Nicol, W. H. Sanders, and K. S. Trivedi, "Model-based evaluation: from dependability to security," *IEEE TDSC*, 2004.
- [19] X. Ou, S. Govindavajhala, and A. W. Appel, "Mulval: A logic-based network security analyzer," in *USENIX Security Symposium*, 2005.
- [20] X. Ou, W. F. Boyer, and M. A. McQueen, "Mulval tool download and readme page," 2006, <http://people.cs.ksu.edu/~xou/argus/software/mulval/readme.html>.
- [21] —, "A scalable approach to attack graph generation," ser. CCS '06. New York, NY, USA: Association for Computing Machinery, 2006.
- [22] M. Pendleton, R. Garcia-Lebron, J.-H. Cho, and S. Xu, "A survey on systems security metrics," *ACM Comput. Surv.*, vol. 49, no. 4, Dec. 2016.
- [23] A. Ramos, M. Lazar, R. H. Filho, and J. J. P. C. Rodrigues, "Model-based quantitative network security metrics: A survey," *IEEE Communications Surveys Tutorials*, vol. 19, no. 4, pp. 2704–2734, 2017.
- [24] K. Scarfone and P. Mell, "An analysis of cvss version 2 vulnerability scoring," in *Proceedings of the 2009 3rd International Symposium on Empirical Software Engineering and Measurement*, ser. ESEM '09. USA: IEEE Computer Society, 2009, p. 516–525.
- [25] A. Walker, "Micro-segmentation network security software," <https://www.ditno.com/micro-segmentation-network-security-software>.
- [26] L. Wang, T. Islam, T. Long, A. Singhal, and S. Jajodia, "An attack graph-based probabilistic security metric," in *DAS XXII*, 2008.
- [27] L. Wang, A. Singhal, and S. Jajodia, "Measuring network security using attack graphs," in *QoP*, 2007.
- [28] A. Wool, "A quantitative study of firewall configuration errors," *Computer*, vol. 37, no. 6, pp. 62–67, 2004.
- [29] Yang Wang, D. Chakrabarti, Chenxi Wang, and C. Faloutsos, "Epidemic spreading in real networks: an eigenvalue viewpoint," in *ISRDS*, 2003.
- [30] S. Zhang, X. Ou, A. Singhal, and J. Homer, "An empirical study of a vulnerability metric aggregation method," 01 2011.