SHIELDX

WHITE PAPER

Machine Learning for Effective Micro-Segmentation

According to the Global Information Security Workforce Study from (ISC)2, the worker shortage in cybersecurity will rise to 1.8 million by 2022, increasing by 20% from 2015. But the gap in skilled workers is already evident, as "66% of respondents reported not having enough workers to address current threats"[1].

TIME IS OF THE ESSENCE

The widespread adoption and complex mixes of highly-virtualized data centers, private and public cloud deployments—also referred to as multi-cloud environments—have forced a paradigm shift upon IT organizations. Their primarily perimeter-based security strategies are no longer equipped to protect against evolving threats and attack techniques, especially their lateral movement and progressive access to highly valuable assets and subsequent exfiltration of sensitive data. As a result, micro-segmentation has emerged as a leading strategy for securing these environments.

FIGURE 1: WORKER SHORTAGE IN CYBER SECURITY

Unfortunately, most organizations today must consider whether they have the resources necessary to implement micro-segmentation with its extensive up-front work of documenting security requirements and formulating policies. And of those resources, time is the most valuable asset for security professionals, and the most scarce.

To help IT organizations meet this challenge, security vendors should enhance solutions with new technologies and automation capabilities to relieve security professionals of the manual tasks of micro-segmentation, accelerating deployment and giving them more time to use their skills to address imminent threats.

FROM A FLAT TO A MICRO-SEGMENTED NETWORK

Organizations that understand the need for micro-segmentation and decide to undertake micro-segmentation projects find themselves confronted with big questions and hurdles: How do we get from a flat or lightly-segmented multi-cloud to a micro-segmented environment? How do we group workloads for the most efficient application of policy? What should that policy be? And how much time will this really take?

For most organizations the hard truth is: this is done manually. It is an enormous hindrance that negatively affects the adoption of micro-segmentation. To understand the extent of the upfront investment it takes to implement micro-segmentation, let's look at a simple example.

Company X has 15,000 workloads across 500 applications in their multi-cloud environment. For this calculation, we assume that their applications are built using an n-tier (also multi-tier) architecture. An n-tier architecture organizes software components into tiers that each provide dedicated functionality. Figure 2 shows the most common implementation of the n-tier architecture which is built with three tiers: a client (e.g., web) tier that implements the user interface, an appli-

(1) International Info System Security Certification Consortium (ISC)2 (2017). Global Cybersecurity Workforce Shortage to Reach 1.8 Million as Threats Loom Larger and Stakes Rise Higher. https://www.isc2.org/en/News-and-Events/Press-Room/Posts/2017/06/07/2017-06-07-Workforce-Shortage.

cation tier which realizes the business logic for the application, and a data management tier that stores information used by the application.

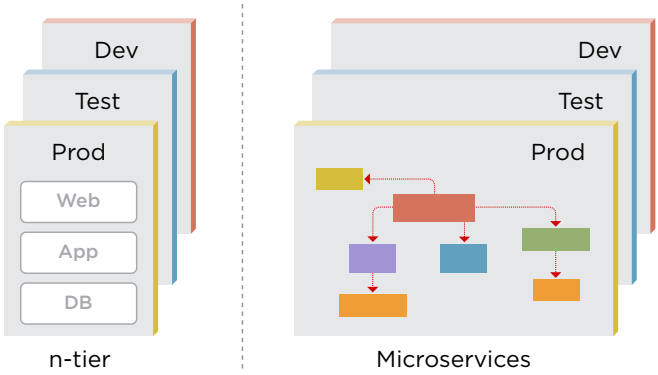

FIGURE 2: APPLICATION ARCHITECTURES

The modularized design of this n-tier architecture allows for easier development, testing, and maintenance of the application. While this is the most commonly found architecture today, companies with other application architectures, like a microservices-based architecture (Figure 2), face the same issue when grouping workloads and micro-segmenting a network.

Going back to the example, a three-tier architecture means that 500 applications translate into as many as 1500 tiers. In a typical multi-cloud environment, three instances of an application are provisioned, one each in Development, Test and Production environments. This amounts to a total of up to 4500 tiers.

Micro-segmentation generally requires policy assignment at the level of the tier, but policy enforcement at the level of the individual workload. A policy set comprises a combination of access control, threat detection, and malware policy. Manually placing 15,000 workloads, each of which is part of at least one application or service, across 4500 tiers and assigning a policy set to each of those tiers is an enormous task and demands considerable time-investment from security professionals that already don't have enough time.

To put a number to this investment, let's assume that a security professional would need one minute per each of the 15,000 workloads for research, and to place them in a group and assign the correct policy set. In a perfect world, this would amount to nearly six person-weeks of work. In reality, however, records of these workloads are not readily available as applications are usually owned by several different teams across enterprises. Hence, it is very likely that this undertaking would consume much more time than this conservative estimate.

MACHINE LEARNING FOR THE AUTOMATION OF MICRO-SEGMENTATION

The solution to this problem lies in the automation of the process. A network-based security solution that attempts effective micro-segmentation must automate:

- network discovery
- workload grouping and policy assignment,
- and visualization of the micro-segmented network.

This automation can be accomplished with the help of machine learning. Automating the grouping of workloads will save time as it relieves security professionals from manually collecting connection records and at the same time reduces complexity. This will often mean the difference between a successful and a failed micro-segmentation project.

Network discovery

Constantly migrating workloads and changing networks—the essence of multi-cloud environments—cause organizations headaches when keeping track of their data center topology. An ideal micro-segmentation solution will help them discover workloads and identify patterns of connection amongst them as part of the micro-segmentation process. Data used for this initial discovery could be a record of all flows within the environment over a defined period of time, which can be used as a starting point for the approximation of least privilege. But how is this data obtained?

(2) Liu, Ling, Özsu, M. Tamer (Eds.) (2009). Encyclopedia of Database Systems, pp 1862-1865.

The discovery process should start with a pervasive and non-intrusive deployment of a micro-segmentation solution in a multi-cloud environment. The solution may be deployed in tap mode, inspecting the network traffic and collecting flow records. It is in the best interest of any organization attempting effective micro-segmentation that this initial deployment monitor as close to every workload as possible to gather comprehensive connection data.

A way to use the obtained connection records for application classification is deep packet inspection. This will allow the solution to determine the appropriate Threat Prevention, Malware Detection and Data Loss Prevention policies.

Organizations should be able to ensure the data collected matches what they have running in their multi-cloud environments. As data center topologies are best consumed by humans in a visual fashion, connection and application detection events should be aggregated to build a visual representation of the workloads, the groups they belong to, and their interactions.

Workload grouping and policy assignment

To be part of the same group, workloads must share identical security requirements. These requirements can be determined by observing the following workload characteristics:

1. Workload behavior
2. Workload software configuration
3. Workload function

For applications built with an n-tier architecture, the discovered groups will correspond to the tiers of the application. For other architectures, such as a microservices-based application, the groups will simply represent sets of workloads with identical security requirements.

As a result of the initial discovery process, a micro-segmentation solution using machine learning could recommend a scheme for the grouping of workloads and the appropriate policies to apply to those groups. However, it is important that these recommendations be reviewed by a security professional before implementing them. Having this validation process is especially important for two reasons:

1. Misconfigured or already-compromised workloads may be making connections that shouldn't be allowed.
2. Connections that need to be permitted might not occur during the data collection period.

The human element cannot (and should not) be fully eliminated from the process. Security professionals know their environments and can quickly make adjustments and fine-tune as needed.

Visualization of the micro-segmented network

To present the micro-segmented network in a fashion familiar to the user, discovered groups should be assembled into application instances which reside in multiple environments, such as Development, Test and Production (Figure 3), as part of the automation process. In the case of a n-tier architecture, where groups map into tiers, an instance encompasses n tiers.

During the discovery process, machine learning should not only determine groups of workloads, but also the connectivity patterns between them. For instance, using the set of rich data points from the discovery process can help to identify sets of instances that make up a complete application. In the same manner as during discovery, the grouping process

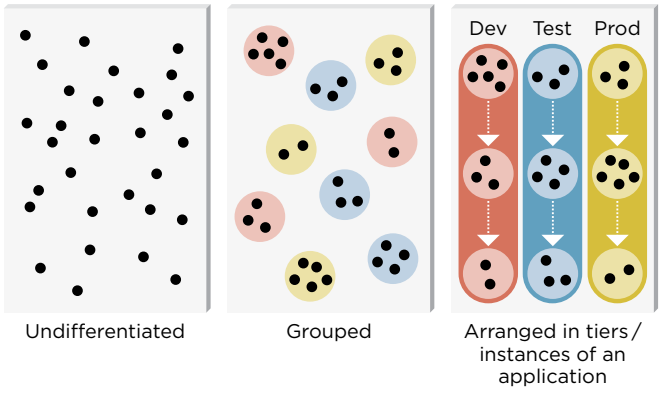

FIGURE 3: ASSEMBLING APPLICATION INSTANCES IN MULTIPLE ENVIRONMENTS

should allow for assembled application instances to be inspected, and the process tuned to allow for adjustments, if needed.

To effectively help organizations carry out micro-segmentation projects, the ideal solution should present an accurate map of all discovered workloads, application instances, and their inter- and intra-connectivity in a fashion that is easy to comprehend visually. It is also important that the visual representation always adapts and reflects the particular application architecture (n-tier, microservices, service-based, etc.) on which the application is based.

There is substantial benefit from this process for a solution that is network-based, as opposed to one that is host-based. Modifying each workload when implementing a host-based solution means that the impact of agent installation on the workload behavior and performance for all applications would have to be tested before the discovery process could take place. The testing process required before the deployment of a network-based solution in a non-intrusive discovery mode is much shorter and less complex. For a large deployment, it is likely that with a network-based approach, the necessary testing followed by the discovery process and product deployment could be completed before the testing of the host-based solution is done.

SUMMARY

Micro-segmentation adds another layer of security needed to protect against today's evolving threat landscape. But, to date the implementation of micro-segmentation consumes enormous resources, forcing organizations to make trade-offs. Micro-segmentation projects will only become feasible by incorporating machine learning strategies.

When evaluating micro-segmentation solutions, machine-learning-enabled automation should be considered as a critical component of the solution, rather than just an enhancement. The costs that organizations are facing when forced to manually implement micro-segmentation are inhibiting the adoption of the technology and contributing to a reduction in security. Given the shortage of security professionals and the value of their time, organizations can significantly improve their security and decrease their total cost of ownership with a micro-segmentation solution that employs machine learning.

REQUIREMENTS FOR A MICRO-SEGMENTATION SOLUTION

- ☐ Discovery of existing data center assets such as workloads, networks, etc.
- ☐ Automation of workload grouping and policy assignment
- ☐ Mapping data center topology with multiple application-centric visualization choices

About ShieldX

ShieldX is redefining cloud security to better protect organizations against cyber threats—regardless of where sensitive data resides or how it moves across public, private, or multi-cloud environments. Organizations are using APEIRO to scale security and micro-segmentation on demand, support business innovation, meet compliance requirements and protect against the latest cyberattacks. Based in San Jose, CA, ShieldX was founded in 2015 and is privately funded.

SHIELDX

ShieldX Networks, Inc.
2025 Gateway Place, Suite 400
San Jose, CA 95110 USA

+1 408.758.9400
info@shieldx.com
www.shieldx.com

January 2018