SHIELDX

WHITE PAPER

Security Controls for Effective Micro-Segmentation

According to Ponemon, organizations face an average cost of $4 million per breach[1]. Attackers accomplish this damage through a strategy of rapid and constant innovation of sophisticated hacking techniques to make it impossible for organizations to foresee and protect against all possible attacks, leaving room for their exploits to succeed. Attackers also use the advantage of an "inherently low cost of entry to perform attacks"[2] making it only more profitable for them when they do eventually achieve their goals.

SECURITY IS ADAPTING, BECAUSE IT MUST

Most of today's critical business infrastructure has evolved into collections and complex mixes of highly-virtualized data centers, private and/or public cloud deployments, and provider services. Or simply, a 'multi-cloud' environment.

IT security strategies, tactics and teams must continuously adapt to keep pace with these fast-paced innovations in application and infrastructure technology, as well as the ever-changing techniques employed by attackers. Unfortunately, while there have been tremendous advances and accomplishments in the fight to keep organizations and their data safe from attacks, enterprises worldwide are breached constantly.

As a result, it has become widely accepted that the traditional, perimeter-based approach to securing these multi-cloud environments is insufficient for addressing the full range of today's highly complex attack activities. Placing network security only at the ingress point of multi-cloud environments leaves security professionals unable to stop an attack after the initial breach, maintain visibility or track its progression.

With the rise of software-defined networking, Forrester's Zero Trust Model[3] and its application within multi-cloud environments has become a leading strategy to help protect against the lateral movement of an attack, its access to high-value assets, and its subsequent exfiltration of sensitive data. One of the foundational, supporting elements of this strategy is the use of micro-segmentation. But while ESG research shows that "68% of enterprise organizations use some type of software-based micro-segmentation technology"[4], its implementation still presents some challenges and fails to protect against today's most pervasive assaults.

ACL-BASED MICRO-SEGMENTATION—NOT ENOUGH

Micro-segmentation is the software-based extension of the traditional practice of segmenting a network via firewalls, subnets, VLANs, etc. It breaks down a network into individual workloads to enforce fine-grained access control between those workloads. In a micro-segmented network, security perimeters are no longer limited to the ingress point of a data center and a few coarse-grained security zones. Instead, perimeters are established around each workload within a multi-cloud environment.

---

(1) **Ponemon Institute. (2016).** 2016 Cost of Data Breach Study: Global Analysis Retrieved from https://securityintelligence.com/media/2016-cost-data-breach-study/

(2) **Jacobson, D., & Idziorek, J. (2013).** Computer Security Literacy. Staying Safe in a Digital World. Boca Raton, FL: Taylor & Francis Group, LLC.

(3) In 2009, Forrester developed a new information security model, called the Zero Trust Model, which has gained widespread acceptance and adoption.

(4) **Oltsik, Jon (2017).** Micro-segmentation projects span enterprise organizations: http://www.csoonline.com/article/3187176/security/micro-segmentation-projects-span-enterprise-organizations.html

Micro-segmentation is based on the *Principle of Least Privilege*, and is used to disrupt post-breach activity. The Principle of Least Privilege establishes that in any abstraction layer of a computing environment, every module (such as a process, a user, or a program, depending on the subject) must be able to access only the information and resources necessary for its legitimate purpose. In the multi-cloud application, this means that each workload should only be permitted to make connections that are required and necessary to accomplish its tasks, without exception.

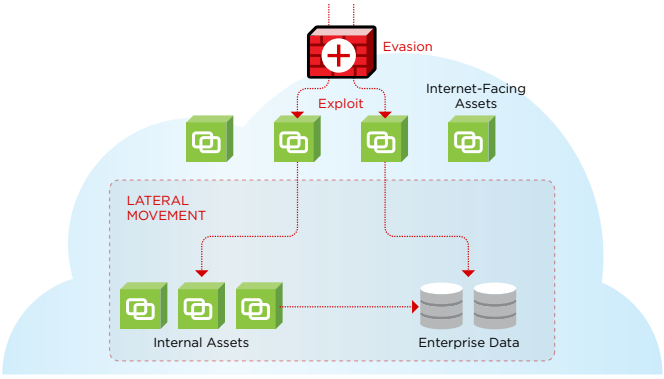

FIGURE 1: PERIMETER-BASED SECURITY

To understand how micro-segmentation helps secure networks, let's look at the anatomy of an advanced attack.

Figure 1 illustrates a security approach that relies exclusively on perimeter-based protection at the ingress point. While the odds of any individual attack evading these protections are quite low, a typical attacker will boost these odds by generating a very high volume of attacks. This is achieved through automation which also enables a marginal cost per attack that approaches zero. As a result, attackers will evade the perimeter-based protection and exploit an internet-facing asset.

Although best practices dictate that sensitive data not be stored on internet-facing assets, attackers follow the initial evasion and begin a process of moving laterally through the environment to progress toward a high-value asset. With no defenses behind the perimeter protection, a successful initial breach translates into unrestricted movement for the attacker, allowing the exploitation of internal assets and achievement of the final goal—the exfiltration of sensitive enterprise data.

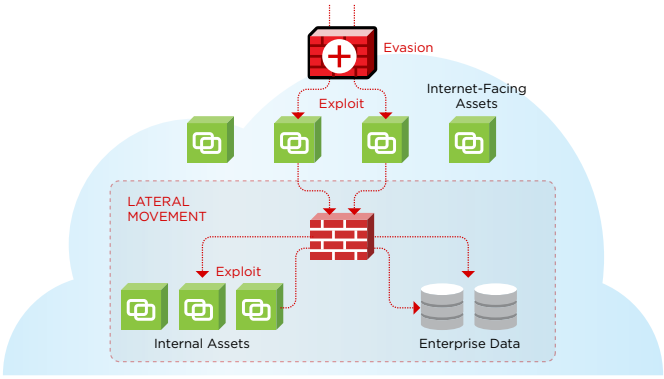

FIGURE 2: MICRO-SEGMENTED NETWORK

When micro-segmentation is applied to the same configuration as *Figure 1*, another layer of protection is added, as depicted in *Figure 2*. Micro-segmentation, as typically implemented, provides protection through basic ACLs (access control lists) that allow only permitted connections to be made within the environment. Therefore, following the initial breach, the attacker is no longer able to freely move within the multi-cloud environment. Instead, if a compromised internet-facing asset attempts a connection (destination, protocol, and port) that isn't required for normal operation, that connection is denied.

However, in practicality, ACLs may slow down a persistent attacker, but they won't stop the attacker entirely. As there are many types of connections that must be made within any multi-cloud environment, an attacker can take advantage of this by breaching an asset, and then moving laterally via connections and protocols that are normally permitted for the breached victim. In this case, the attacker enlists the assistance of the breached machine in a compromise formally known as a *Confused Deputy.* The deputy is 'duped' into using its assigned privileges to move toward the high-value target. In this manner, the attacker can move laterally without detection, even with ACLs in place.

The problem of the Confused Deputy is not a new one, nor is the need for security beyond simple ACLs. In the late 1980s, firewalls designed to impose basic ACL policy became commonplace for internet-connected organizations. Threats like the Morris worm caused many organizations to deploy a firewall to enforce the Principle of Least Privilege at the perimeter.

But IT teams soon realized that additional security controls were needed to limit the range of successful attacks and protect the network. They extended the Principle of Least Privilege from Layers 3/4 to Layer 7 through advanced, DPI-powered (deep packet inspection) security controls, including intrusion detection and prevention, malware detection, data loss prevention, anomaly detection and reputation analysis *(Figure 3).* Today, we couldn't imagine perimeter security without these advanced controls.

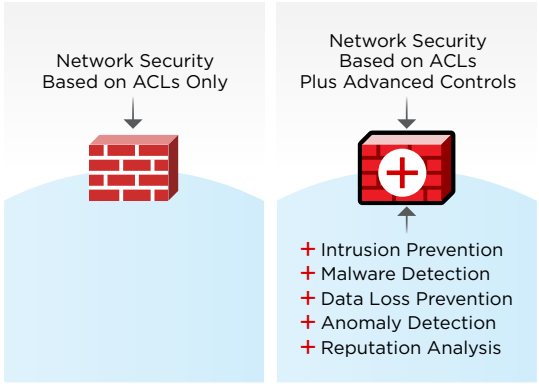

FIGURE 3: ADVANCED SECURITY CONTROLS

Although micro-segmentation is a necessary first step towards protecting against advanced threats and lateral movement, alone it is simply not enough, since an attacker can use the privileges of Confused Deputies to move laterally and to exfiltrate sensitive data. Micro-segmentation is an important foundational element for building a multi-layered security framework in virtualized and orchestrated environments, but it is by no means sufficient by itself.

SECURITY CONTROLS FOR COMPREHENSIVE LATERAL PROTECTION

Spanish-American philosopher George Santayana once wrote, "Those who do not remember the past are condemned to repeat it." Just as we learned at the perimeter 25 years ago, using ACLs alone to protect the East/West axis of a network won't get the job done. Micro-segmentation that employs enterprise-class, DPI-based security controls is required, not only on the perimeter but also within multi-cloud environments.

To fully protect a network and solve the Confused Deputy problem, the Principle of Least Privilege must be extended from L3/4 to L7 *(Figure 4).* Applying least privilege at L3/4 is relatively simple and achieved through standard ACLs—whitelisting. L7 protocols, however, are much more numerous and complex, which makes an approach based on whitelisting impractical. Instead, blacklisting can be employed. Creating a policy that defines blacklisted behaviors and activities at L7 and actively enforcing that policy solves the Confused Deputy problem and completes the job that is only partially addressed by ACLs alone.

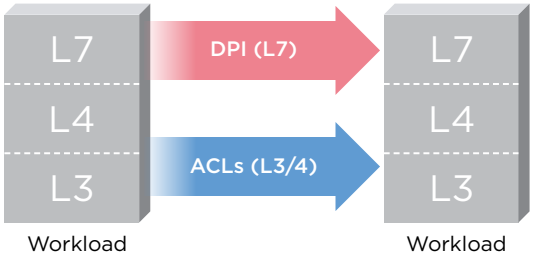

FIGURE 4: EXTENDING LEAST PRIVILEGE TO L7

Since individual security controls tend to target different behaviors and activities, blacklisting is most effective when the broadest possible range of these controls can be applied. Hence, a proper solution for extending least privilege to L7 should include a full range of such controls, such as ACLs with AppID, intrusion prevention, malware detection, data loss prevention, anomaly detection (including both flow meta data and L7 content) as well as TLS decryption, which ensures visibility at L7.

A new model for comprehensive, network-based security that uses the latest innovations in application and infrastructure technology has to extend least privilege to L7. This extension of micro-segmentation must include a full set of security controls that provide network protection as well as visibility, automation, and coordination across multi-cloud environments. Consequently, the breach of an asset in one area of the network can be detected, reported, and contained and will not spread and compromise others.

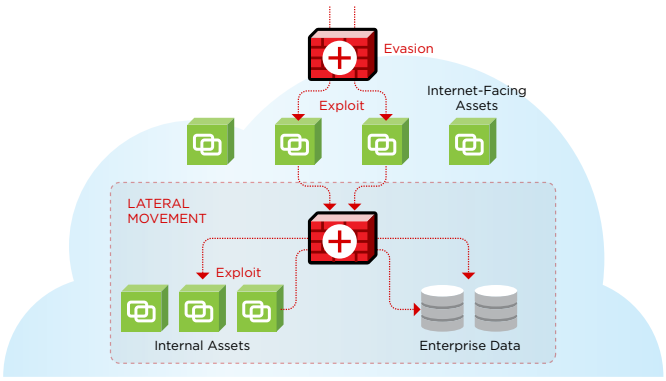

FIGURE 5: EFFECTIVE MICRO-SEGMENTATION

Building upon our previous example, *Figure 5* shows how effective micro-segmentation protects multi-cloud environments. ACLs, coupled with DPI-powered security controls that inspect lateral traffic permitted by those ACLs, add the security necessary to prevent attackers who have already breached the perimeter from moving laterally. High-value assets and the sensitive data they contain remain secure.

The continuous success of attacks based on exploits, such as EternalBlue or EternalRomance, illustrates how effective micro-segmentation can protect against these types of lateral attacks. They evade protection at the perimeter and then propagate laterally using common protocols, such as Server Message Block version 1 (SMBv1). ACL-based micro-segmentation, even if properly configured, could only slow their propagation as connections like SMBv1 are commonplace and must be permitted within the typical multi-cloud environment.

A proper solution should inspect traffic fully, inline at the packet, flow and file level and detect a malicious connection from the Confused Deputy. It should use Threat Prevention to stop the attack in real-time, and actively prevent an attack from spreading and eventually, compromising the organization and its assets.

SUMMARY

While solutions available today offer simple ACL-based micro-segmention for the multi-cloud, none can fully secure these environments at all layers of the protocol stack in a way that is also cost-effective and practical to operationalize.

When evaluating new security solutions for multi-cloud environments, IT organizations should set the criteria to require micro-segmentation extension to L7, with a full set of DPI-powered network security controls, in a format that can be effectively deployed. Then, even if an initial breach occurs, their organization will be able to detect an attack's complex progress through their environment over long periods of time, in multiple stages and choose to take action, before damage is done.

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

September 2017