"Digital Transformation: Trust but Verify"
Exclusive to *Network Computing*

Moving digital assets to the public cloud reduces costs and increases productivity, but it poses some new information security challenges. Specifically, many Intrusion Detection and Prevention systems (IDPS) that were designed for the on-premises network come up short when deployed in the public cloud. For this reason, public cloud providers have built-in security layers to manage information security using their own security monitoring infrastructure. Unfortunately, these built-in monitoring services are one-size-fits-all and may miss crucial customer-specific security requirements or user account compromises. This leaves cloud-based assets more vulnerable to data breaches. In this article, we'll look at the challenges of securing cloud-based assets.

Why public clouds are difficult to secure

Public clouds are great when it comes to providing shared compute resources that can be set up or torn down quickly. The cloud provider offers a basic software interface to provisioning storage, servers and applications, and basic security monitoring that runs on top of that interface at the application layer. But the application layer runs on top of the network, and the network is the only place where certain classes of dangerous security breaches can be detected and prevented.

In the cloud, customers can't conduct network-level traffic analysis, because public clouds don't give customers access to the network layer. Clouds restrict users from inspecting or logging the bits that go over the network wire. Inspecting a public cloud at the application layer can give customers information about what the network endpoints are doing, but that's only part of the picture. For example, breaches due to users' misbehavior are only visible at the network layer by observing the communication patterns that are inconsistent with company policies. The cloud's built-in monitoring services would not be aware of it because they do not monitor network behavior on behalf of the enterprise. Importantly, if malware or a rogue application somehow makes it into a cloud instance or remote VM hosted in the cloud, native cloud monitoring services may not detect its malicious behavior at the network level. Because customers don't have access to the bits being transmitted, they'll never know the malware is there.

And the network threats are there. An April 3, 2019 article in ZDNet mentioned that over 540 million Facebook records were exposed on AWS. In 2017, 57 million Uber customer records were compromised because hackers extracted Uber's AWS credentials

from the company's private GetHub account. Public cloud offers no tools for monitoring the network data that would have detected and prevented these breaches.

Public cloud operators could see what's going on if they were to look at the network traffic, but they don't provide that information to their customers. Most of the time, public cloud operators are focused on providing application-level security information from systems like firewalls or endpoint Antivirus solutions. Adding NG firewalls from third-party vendors to public cloud deployments adds the ability to customize the inspection of all the bits flying by, but fails to detect communications within the cloud (for example, between a web server and a database) or lateral communications (for example, a compromised host trying to spread within the internal cloud network between VMs). This leaves blind spots that can allow malware to execute without the user's knowledge. Lastly, when there is a breach, in most cases, cloud customers can't even quantify, precisely, the number of records or the amount of data exfiltrated.

As it's not feasible to deploy hardware on a public cloud provider's premises, the way to eliminate these blind spots lies with software that can implement a virtual tap and monitor traffic at the network level. The security vendor industry is now moving away from dedicated hardware devices and toward multi-function software that will address these needs.