11 LIMITATIONS OF NETWORK-CENTRIC SECURITY IN THE CLOUD

WHY YOUR OLD APPROACH MAY PUT YOU AT RISK

Lacework™

INTRODUCTION

Traditional security solutions were designed to identify threats at the perimeter of the enterprise, which was primarily defined by the network. Whether called firewall, intrusion detection system, or intrusion prevention system, these tools delivered “network-centric” solutions. Innovation was slow because activity was dictated, in large part, by the capabilities, and limitations, of available technology resources. Much like a sentry guarding the castle, they emphasized identification and were not meant to investigate activity that might have gotten past their surveillance.

Modern threats targeting public clouds (PaaS or IaaS platforms) demand a different level of insight and action. They operate differently than traditional datacenters: executables come and go instantaneously, network addresses and ports are recycled seemingly at random, and even the fundamental way traffic flows have changed. To operate successfully in modern IT infrastructures, you have to reset how you think about security in cloud.

Surprisingly, many organizations continue to use network-based security and rely on available network traffic data as their security approach. It's important for decision makers to understand the limitations inherent in this kind of approach so they don't operate on a false sense of security.

> "**92%** of enterprises tried to secure cloud workloads using traditional network security controls. **74%** had to abandon some or all of these controls because they proved to be incongruous with this new use case." – *Enterprise Strategy Group*

To help security professionals understand the new world of security in the cloud, below are 11 specific use cases where network-centric security is inadequate to handle the challenges of security in modern cloud environments.

Read more: What's Holding Back Enterprise Security Technology Transformation, ESG, Oct 2017.

OVERVIEW

1. NETWORK-BASED DETECTION CREATES MANY FALSE POSITIVES
2. CYBER ATTACKS HAVE MOVED BEYOND THE NETWORK ATTACK SURFACE
3. NETWORK DATA CAN'T ATTRIBUTE CLOUD SESSIONS TO ACTUAL USERS
4. DEEP PACKET INSPECTION ON NETWORK TRAFFIC DOESN'T SCALE
5. MALICIOUS ACTIVITY AT THE STORAGE LAYER IS NOT DETECTED
6. NETWORK-BASED SECURITY IS BLIND TO CONTAINER TRAFFIC
7. CONTAINER ORCHESTRATION TRAFFIC IS INCOMPLETE
8. FILE INTEGRITY MONITORING IS NOT COVERED BY NETWORK LOGS
9. FILE-BASED MALWARE DETECTION IS NOT POSSIBLE
10. CAPTURING ALL NETWORK TRAFFIC IS NOT FEASIBLE
11. CAN'T SEE UNINTENTIONAL MISCONFIGURATIONS IN NETWORK LOGS

1 NETWORK-BASED DETECTION CREATES MANY FALSE POSITIVES

Nothing has confounded network security as much as the demise of static IP addresses and endpoints in the cloud. Endpoints used to be physical; now they are virtual and exist as containers. In the cloud, everything is dynamic and transient; nothing is persistent. IP addresses and port numbers are recycled rapidly and continuously, making it impossible to identify and track over time which application generated a connection just by looking at network logs. Attempting to detect risks, and threats using network activity creates too many irrelevant alerts and false positives.

"Many SOCs within large enterprises are challenged to respond to 5-10% of received alerts." — 451 Research, Security Automation and Orchestration Bring Sanity to Incident Response Chaos, May 2018

451 Research™

2

CYBERATTACKS HAVE MOVED BEYOND THE NETWORK

Illicit activities have moved beyond the network attack surface in the cloud. Here are four common attack scenarios that involve configuration and workloads (VMs or containers) in public clouds, but will not appear in network logs:

1. User privilege changes: most cyber attacks have to operate a change of privilege to succeed.
2. The launch of a new application or a change to a launch package.
3. Changes in application launch sequences.
4. Changes made to configuration files.

3

NETWORK DATA CAN'T ATTRIBUTE CLOUD SESSIONS TO ACTUAL USERS

The common DevOps practice of using service and root accounts has been a double-edged sword. On one hand, it removes administrative roadblocks for developers and accelerates even further the pace of software delivery in cloud environments. On the other hand, it also makes it easier to initiate attacks from these "privileged" accounts and gives attackers another place to hide.

By co-opting a user or service account, cybercriminals can evade identity-aware network defenses. Even correlating traffic with Active Directory can fail to provide insights into the true user. The only way to get to the true user of an application is to correlate and stitch SSH sessions, which is simply not possible with network only information.

4 DEEP PACKET INSPECTION ON NETWORK TRAFFIC DOESN'T SCALE

When security based on simple network characteristics (addresses and port numbers) started to fail at stopping increasingly sophisticated attacks, the security industry responded with DPI (Deep Packet Inspection) and "next-gen firewalls" to bring a different level of insights at the application and user level. These still add value in the cloud when customers use packaged applications (SaaS), but with IaaS and PaaS platforms, applications are custom-built. Because next-gen firewalls can't understand custom applications (at least not without a herculean effort), they can't detect attacks.

IT leaders must also consider the growing use of host encryption. Trying to use standard network tools at cloud-scale makes packet inspection extremely difficult.

5

MALICIOUS ACTIVITY AT THE STORAGE LAYER IS NOT DETECTED

In cloud environments, the separation of compute and storage resources into two layers creates new direct paths to the data. If the storage layer is not configured properly, hackers can target APIs and conduct successful attacks without being detected by network-based security.

On AWS specifically, S3 bucket misconfigurations common and have left large volumes of data exposed. Data leaks due to open buckets will not appear on network logs unless you have more granular information that can detect that abnormal activity is taking place.

6

NETWORK-BASED SECURITY IS BLIND TO CONTAINER TRAFFIC

Network logs capture network activities from one endpoint (physical or virtual server, VM, user, or generically an “instance”) to another along with many attributes of the communication. Network logs have no visibility inside an instance. In a typical modern micro-services architecture, multiple containers will run inside the same instance and their communication will not show up on any network logs. The same applies to all traffic within a workload.

Containerized clouds are where cryptocurrency mining attacks often start, and network-based security has no ability to detect the intrusion.

7

CONTAINER ORCHESTRATION TRAFFIC IS INCOMPLETE

Most containerized cloud deployments use orchestration tools such as Kubernetes to manage the lifecycle of their containers at scale. You must secure orchestration tools and traffic between containers and orchestrations tools. Attributing the traffic to the correct container requires visibility in the corresponding name space – this is absent from network logs.

These “holes” within and around containers create corners where hackers can plant bad code or siphon off private data and are undetectable by network security tools.

Read more: Containers At-Risk: A Review of 21,000 Cloud Environments

8

FILE INTEGRITY MONITORING IS NOT COVERED BY NETWORK LOGS

Many regulations (HIPAA, PCI, FISMA, NERC-CIP) require File Integrity Monitoring (FIM) because file changes are leading Indicators of Compromise (IoC). Hackers will often change an executable, change config files, or delete log files. Network logs cannot provide any information about file level changes and cannot contribute to threat detection or to the compliance audit process.

Visibility into file changes can surface anomalies early especially when combined with the detection of anomalies in application behavior.

9

FILE-BASED MALWARE DETECTION IS NOT POSSIBLE

One of the most important indicators of compromise (IoC) is still file hashes. Known vulnerabilities identified through file hashes should be addressed in priority. Network logs, however, are of no help in this regard because they have no information on file hashes or packages.

10

CAPTURING ALL NETWORK TRAFFIC IS NOT FEASIBLE

Traffic inside the cloud (aka East-West traffic) is up to five times more voluminous than traffic between clients and servers (aka North-South traffic). The associated volume of network logs becomes overwhelming and costly fast if logs need to be stored and analyzed over a long period of time. Even worse, storing information that tracks a five-tuple talking to another five-tuple is useless, especially if it still does not provide the information needed for accurate threat detection and investigations.

Mapping how applications communicate with each other provides more insights. In the cloud, it does not matter if it happened on one machine, one IP, one port, or thousands of different ones. A better performing, less costly approach is to take a logical view of cloud activity at the application level.

11

NETWORK LOGS DO NOT DETECT MISCONFIGURATIONS

Many hackers are continuously scanning networks for interesting assets left unprotected. Network traffic will be alerting on these attacks even though the majority of them are false positives. The ideal solution is to identify the front-facing applications and only alert if an attack happens to an application which is not front facing.

RECOMMENDATIONS

Our recommendations for a secure cloud environment:

1. Get visibility beyond network traffic
2. Embrace solutions that deliver end-to-end visibility
3. Get visibility across cloud accounts
4. Look for solutions that cover existing, new and upcoming cloud architecture – virtualized, containerized, and serverless
5. Develop a baseline for normal behavior, and identify behavior that presents a threat
6. Better analysis provides fewer, but more actionable, alerts

CONCLUSION

Focusing exclusively on network connections is not enough to secure cloud environments. Servers and endpoints don't yield any better results as they come and go too fast for an endpoint-only strategy to succeed. So, what can you do?

At Lacework, we decided to take a different approach. We collect data at the VM and container level, organize that data into logical units that can give us security insights, and then analyze the situation in real-time. In other words, we go deep vertically when collecting data from workloads, but analyze the information horizontally across your entire cloud. This is how we focus on the application's behaviors and not on network five tuples or single machines.

Interested in more?
Try Lacework for free & validate the security of your cloud:

Streamline security for AWS, Azure, and GCP. Gain unmatched visibility, ensure compliance, and enable actionable threat intelligence.

www.lacework.com/free

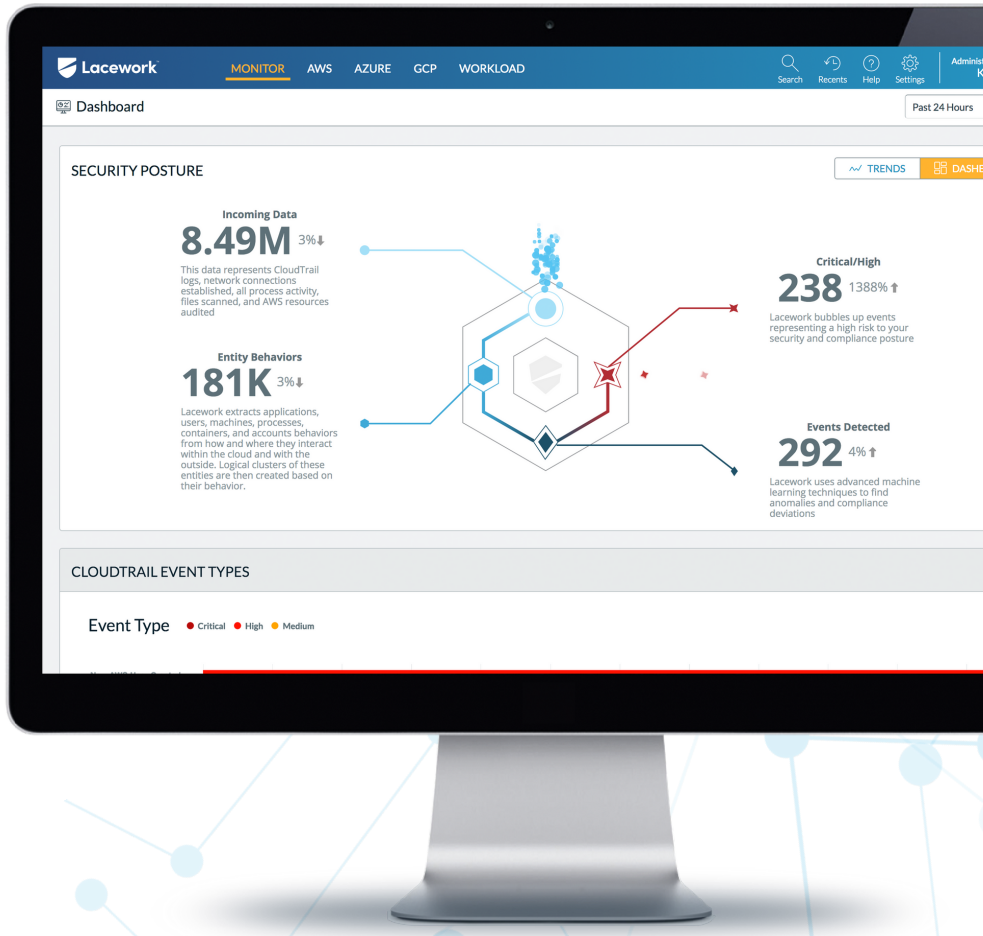

Lacework™