Enterprise Security with a Fluid Perimeter

Matt Bromiley

SANS

A SANS Whitepaper

Enterprise Security with a Fluid Perimeter

Written by **Matt Bromiley**

January 2019

Sponsored by:

Aruba,
A Hewlett Packard Enterprise Company

Introduction

Every time you turn on the news, walk into a store or attend a technical conference, one thing is certain: It's only a matter of time before technology turns even the most far-fetched idea into a reality. User devices grow in complexity and capability extremely rapidly—almost daily, it seems—turning even the most basic of functions into something requiring network connectivity. Meanwhile, enterprises have to bring newer technology on-premises to meet the needs of a demanding workforce. And users keep looking for devices or applications that will create a more harmonious blending of their personal and work lives—and unfortunately, they turn right around and bring these devices and apps into their organizations' networks.

Security organizations have implicated the introduction of user-controlled devices in the corporate network (think BYOD), coupled with cloud-based services and a mobile workforce, in what they refer to as the "vanishing perimeter." At a high level, the idea is simple: In the past, physical network devices were used to control connectivity inside managed networks. User devices and "things that could connect" were limited in scope, distributed by the IT department and monitored for changes. The perimeter was easy to define, and thus easy to secure.

Obviously, most enterprise networks today do not resemble the "old way." Users are more in control of their own devices than ever. Nearly everything can be done with a mobile phone or an app on a tablet—not company-provided. Conference rooms are populated with smart TVs and break rooms with internet-connected vending machines.

Each one brings not only convenience, but also a new potential attack surface. While it's true the perimeter has changed, it has *not* vanished. Instead, the perimeter has become *fluid*—changing daily to adapt to the needs of users and modern business, as opposed to the old way of locking down the enterprise.

Despite this accelerating digital transformation, one thing has remained somewhat, and alarmingly, consistent: the difficulty in finding precisely the right approach to securing the organization, even when armed with knowledge of the perimeter's ever-changing nature. In this paper, we will examine these issues and discuss how your organization can begin to *benefit* from a fluid perimeter. In particular, we will explore:

- The security challenges facing organizations today
- How your organization can identify and harness its fluid perimeter
- Case studies of particular attacker activities, and how they can be detected/prevented with a fluid perimeter

Our goal is to begin changing the way organizations approach security for the devices that enable better user experiences, starting with the realization that your network is likely giving you most, if not all, of the data you need to remain secure. We encourage you to use this paper as a way to assess how your organization is currently dealing with a perimeter that is morphing into a highly adaptable IT environment. After an assessment, please consider whether or not your information security team might update its methodology, tooling or visibility.

Things Aren't What They Used to Be

Before we examine techniques that organizations can use to discover—and benefit from—their fluid perimeter, let's first discuss some of the challenges facing organizations today. It's important to note that, while challenges can cause a change in visibility and/or security posture, they also provide a chance for the security team to increase its situational awareness.

Changes in the organization can often feel like security burdens; instead, turn them into opportunities to increase your situational awareness and visibility.

Internet of Things

It's imperative that we discuss one of the most significant organizational security risks of late: the Internet of Things (IoT). Inherently, IoT devices are not bad. We may argue as to whether IoT devices are *necessary*, but they are typically not malicious by design. The issue with IoT, from an enterprise perspective, is that the devices can bring uncontrolled risks into the environment and thus erode the known perimeter. Additionally, IoT devices are often built on open source software platforms, such as Linux, which the organization may not fully support.

When IoT devices are placed within the environment—whether supported by the organization or not—the perimeter and security posture must adjust to include these attack surface(s). From a security perspective, this may mean placing IoT devices within specific subnets, whitelisting only the ports/IPs needed for communication, and/or disabling internet connectivity entirely if it is not needed.

Mobile Workforce and BYOD

Two of the most significant challenges that IT managers have been dealing with for many years are an increasingly mobile workforce and necessary—but often relaxed—BYOD policies. Some departments are now run by dispersed groups of users working on personally owned devices with questionable, if any, security installed. Of course, the perimeter increases in fluidity with every addition of a mobile worker or a BYOD device. How can the organization come close to ensuring the devices are not bringing threats into the company network? Again, the perimeter must be adjusted accordingly, and the organization must achieve security where possible.

With security in mind, it is essential to hold users to a security policy that protects all assets that connect to a corporate network. Ensuring compliance might involve products such as MDM that allow the enterprise to maintain control over a personal mobile device when necessary. Additionally, users should expect some level of endpoint security to be required on their devices (typically using a cloud-based agent that allows for remote check-in, management and reporting without the need to connect to a domain or VPN).

By implementing remote management tools, the security and operations teams can at least begin to get the upper hand with potential software vulnerabilities that users might be exposed to. Some tools also allow for remote takeover, which, in the event of theft or loss, can be essential in protecting corporate data.

The Cloud/Third-Party Providers

Is it possible to discuss security risks and not include the ominous cloud? Likely not—and unfortunately, despite the ubiquity of cloud computing, many organizations still don't consider the risks that the cloud may pose to their perimeters.

Let's split the cloud into two appropriate categories:

1) Servers/systems hosted within the cloud
2) Third-party providers that utilize cloud hosting for their functionality

Each category offers its own level of complexity in managing a security perimeter.

The cloud multiplies your security concerns: Recognize that your perimeter includes your own servers and those of the vendors you use in the cloud.

Systems Within the Cloud

When an organization decides to move servers or functionality to the cloud, it is often performing a cost- or availability-based assessment. Rarely do organizations move to the cloud strictly for security purposes. However, there are some inherent benefits to hosting within a cloud instead of on-premises. These benefits may include centralized logging and data center-level bandwidth, among many others.

Depending on the functionality of the system, the movement of system(s) to the cloud will certainly result in perimeter expansion. Your security team may not be responsible for the security of the entire cloud, but it should still include the cloud systems within the inventory of managed assets. This includes ensuring that cloud systems, even though hosted elsewhere, are still properly patched and updated and have strong authentication measures for access.

Third-Party Services

The outsourcing of enterprise functionality to a third-party provider is a double-edged sword. The organization no longer has to maintain resources to manage and provide the associated service, but the security perimeter must be expanded considerably. The degree to which your perimeter expands depends on the types of authentication and integration of the third-party service.

If your employees use different (non-domain) credentials to the third-party site, and if no connections are made back to the company, then the perimeter only need expand to ensure that connectivity to the third-party site is secure. On the other hand, if the third party is integrated via Active Directory and maintains an authenticated state into the company, then the organization must consider this as part of its perimeter. Again: **Remaining fluid in the understanding of your perimeter will allow third-party services to weave in and out of your security posture.**

The "New" Organization: A Fluid Perimeter

Given the assumption that traditional defenses are no longer adequate, how can an organization establish stronger security within its environment? The path to success, as we've discussed, is to increase your understanding of your network, which provides the most specific outlines of the fluid perimeter. Let's examine three techniques that will give your organization a starting point in moving to better defenses:

- Moving beyond simple endpoint data
- Working with network data
- Integrating and automating the use of network data to respond to events faster

It's quite possible that your organization already has some portion of these techniques implemented. We encourage you to use the following as barometers for your current security ecosystem and think of ways that your organization can benefit from additional datasets, analytics and/or automation.

Endpoints Are Only Half the Picture

In our experience with the industry, a good number of organizations have made significant investments in endpoint management, monitoring and detection capabilities. This is a good sign, because those tools give enterprises the ability to respond to endpoint events, address user issues and handle security concerns quickly. Endpoint products provide valuable insight into the positive or negative status of your environment and should be considered a valuable part of fluid perimeter security.

That said, even with advanced capabilities, endpoints are only half the picture. Endpoints are constantly talking—internally and externally—but too often organizations fail to observe and analyze the resulting traffic. What good is knowing the language if you never hear the conversation? For this reason, the first technique for harnessing your fluid perimeter **focuses on ensuring that your organization is monitoring both the endpoint and the network.**

One critical step to harnessing your fluid perimeter: Focus on ensuring that your organization is monitoring both the endpoint and the network.

Network monitoring also allows your organization to "manage" the endpoints that arrive without the knowledge or control of IT. In previous sections, we have used examples such as third-party vendor appliances and other internet-enabled devices, such as IoT. "Things" do not log. They carry minimal if any protection, and the only way of knowing they have been compromised is to watch their network traffic for changes. Network monitoring compensates for not having endpoint control on unmanaged devices within your network.

Optimally, network monitoring for behavioral changes involves deep packet inspection. This technique can yield not only a wide variety of activity information (up to and including application-layer visibility), but also a rich forensic record of transaction information that dramatically accelerates incident investigation and response. If packets are not available, network activity can be approximated by sources such as NetFlow, which, if the switch or router is properly configured, uses User Datagram Protocol (UDP) to send information such as source and destination addresses, ports, protocol and type of service. DNS, DHCP, web proxy and even firewall logs can also fill in the picture. By capturing data from devices in key, pivotal locations, you can very quickly build a profile of what your normal network looks like for critical assets. We'll discuss analytical techniques next.

For security teams to be most effective, they should be able to detect threats at multiple layers of traffic. Don't limit visibility to the point where the perimeter and internal traffic look the same–it makes your team's job harder!

Working with Network Data

Let's be clear about one thing: The addition of data to an already overloaded security team typically yields yet more data that no one has time to look at. As you go through and increase your organization's visibility and security awareness, you can't overload the analysts with new data and expect them to cope. You must accompany the collection of network data with analysis techniques so that your collection efforts are not wasted.

One of the first things we encourage organizations do with their newly acquired network monitoring data is to establish behavioral baselines. Baseline questions that your team should focus on answering include:

- What or who are the users, systems and devices in our environment, and what behavior should we expect to see from those entities?
- What are the normal and peak hours of activity?
- What does DNS traffic look like on our network—number of requests, ensured capture, etc.?
- Has network traffic changed significantly in terms of volume, timing or destination?

These are merely sample questions geared toward understanding "normal" within your network. And when you know normal, it becomes easier to find the abnormal. Baselining can provide useful alerts of questionable activity, yet it can be prone to false positives from normal, benign business or human operations. For this reason, to get actionable data out of networks and baselines, one option is to turn to advanced analytics.

Yes, we just dropped a buzzword. But not for the sake of saying it—we're talking about true machine-learning (ML) capabilities that are working to understand the intricacies of your organization and assist you in defining normal and abnormal. Teams that run smaller environments with little cross-platform exposure can easily answer baselining questions. However, larger environments that measure their endpoints in the tens or hundreds of thousands need a little help establishing their baselines.

User and entity behavior analytics (UEBA)—the use of unsupervised machine learning to profile user and entity behavior—is an advanced security solution that can help find trouble within your enterprise. This technique focuses on using network traffic, flows and logs to develop models of typical/expected user behavior and watching for activity that falls outside of those expectations. If you are looking at behavioral analytics solutions, make sure it looks at both network and log information—many focus exclusively on logs, which creates serious blind spots. A clear example may be a service account logging into a system via Remote Desktop Protocol (RDP). This is suspicious enough that your analytical engine may fire an alert, which will warrant additional investigation. Some UEBA solutions also use supervised machine learning to look for specific attack behaviors such as command and control or lateral spread—in essence finding the malicious within abnormal behaviors and significantly reducing false positives.

Supervised and unsupervised machine learning are just two examples of the advanced analytics that can turn network and log data into a wealth of actionable information and enterprise insight.

Integrating and Automating Endpoints and Network Data

With the previous two techniques implemented, you may be at a point where you have visibility and advanced security awareness across the organization. Data is being ingested into analytical platforms, and the baseline of your organization has been established. With a baseline established, your operations and security teams are able to quickly detect attacks and potential incidents. Additionally, IT activity data is being processed through advanced analytical engines that are generating statistics about user and device behavioral patterns, and machine learning is finding even the smallest hiccups. Sounds like a dream, right? Yet there's still room to grow.

Don't let network monitoring and analytics overload your analysts. Automate and integrate where you can, and let the network start to defend itself.

Even with network monitoring and advanced analytics implemented, industry experience shows us that organizations continue to suffer from the same problem that existed prior to leveraging any advanced implementations whatsoever: analyst fatigue. The addition of IT behavioral data and advanced detection analytics can lead to more dashboards, more screens and more alerts for your analysts to deal with. It can almost feel like one step forward, two steps back. For that reason, our third fluid perimeter technique is to **use advanced knowledge to integrate and automate defenses.**

When you have established normal vs. abnormal within your environment, your security team will see high-fidelity alerts. The confidence level is high enough that you can have a system such as network access control (NAC) implement a range of actions in response to an attack, from re-authenticating to bandwidth throttling to quarantining to outright blocking. Selecting a response short of a full block can buy time for the security team as it does more investigation.

Integration and automation are not only applicable to well-crafted signatures developed from years of investigative experience—you can also use simple events to enforce, automate and integrate defenses. Consider, for example, a company wireless network that is built for a particular set of devices. Any device that joins the network without the proper system information (such as a MAC address or a specific system identifier), can automatically be quarantined or blocked given the conclusive nature of the problem.

Dealing with the Unknown

We have examined the complexities that enterprises must cope with, as well as techniques that can be implemented to harness your fluid perimeter. Remember, we are not discussing an overhaul of your security team or strategy. Let's examine how a fluid perimeter can be used to detect common attack scenarios that organizations are facing today.

Case Study 1: An Uncontrolled Attack Surface

The problem: Many enterprises have semi-permanent devices that belong to third-party vendors on their physical or wireless networks. While the enterprise provides the networking for this device, it has no insight into the software, hence the security team cannot secure the device. **How does the enterprise effectively monitor an IoT device that it cannot manage, and how can it detect incidents in which this device is used to maliciously advance an attack?**

The solution: The addition of unmanaged, internet-enabled devices within a corporate network (IoT vending machines, for example) is common these days. Unfortunately, these devices have also fallen victim to various attacks in recent years (think IoT botnets, such as Mirai). When the organization has third-party, unmanaged devices within the network, endpoint visibility goes away. **So it must focus on the network as the primary tool here.**

First, the organization should update its understanding of the network—and thus its understanding of the perimeter—to include these devices. The operations team is not expected to maintain software on a third-party device (or even an internal device that is supporting a smart building, for example), but the security team is expected to provide security around the third-party device's connectivity and behavior. Therefore, the first step is to ensure that the team is aware of the device in the first place.

Second, proper network monitoring and analytics should be wrapped around these devices from the moment they are installed. The organization can collect network traffic that can be processed and summarized via deep packet inspection. If network traffic is not available, then volumes and ports/IPs used, among other metadata, can be provided by NetFlow. With this data in hand, the organization has established a baseline for typical network activity.

Now, while this may be easier said than done, to identify any potential incidents, the organization need only wait for any device to escape its expected baseline. There are two situations where this may occur:

- The third party has changed its supporting infrastructure, resulting in new IP addresses, ports and/or traffic volumes. Such changes are not incident-worthy, because they will contribute to reestablishing the baseline for the device.
- The device has fallen victim to an attack and begins to act erratically.

The first situation obviously yields no concern. The second should certainly kick the security team into action.

Case Study 2: Wolf in Sheep's Clothing

The problem: An attacker has compromised legitimate system administrator credentials and is moving laterally throughout the network. How can the enterprise discern between legitimate and malicious account activity?

The solution: Lateral movement by itself is not malicious. Administrators often need to connect from system to system to administer ad hoc patching or system updates, and help desk associates typically connect to multiple systems throughout the day to perform their functions. However, this traffic is *expected* and thus should be part of your baseline. The help desk likely exists within a specific subnet or system cluster, and you can monitor traffic from the help desk to other systems. The same goes for your domain administrators: They should be moving laterally using specially assigned pivot systems.

Other types of lateral movement within the environment—as detected via network monitoring—should be investigated. Ask your security teams, "Is there ever a need for someone in the Finance Department to be moving laterally to the Engineering Department?" Attackers often use any pivot point they can during initial stages of a breach, and this activity can be easily detected.

We discussed advanced analytical techniques earlier in this paper, and UEBA can certainly be used to detect malicious account activity. When coupled with automated and integrated defense mechanisms, an organization can get to a point of automatically detecting, and potentially blocking, abnormal user behavior.

Case Study 3: Data Out of Place

The problem: An attacker is well into the advanced stages of a breach and has prepared data for exfiltration from the enterprise's network. How can the enterprise use a fluid perimeter to detect this traffic?

The solution: When considering this scenario, it's important to point out that an organization should not be relying on exfiltration detection as an effective means of detecting breaches. Attackers will take a significant number of steps within your environment before reaching the "ready to exfiltrate" stage, and we hope that you caught them during an earlier stage. However, if an organization has allowed an attacker to go unabated, all hope is not lost.

For detection, we return to the concepts of baselining and machine learning. An effective environment baseline will provide normal traffic volumes and movement to the security and operations teams; data exfiltration will likely violate these norms and expectations. Depending on where an attacker's pivot point is located, it is possible that your ML algorithms will also detect the unexpected fluctuation in traffic and that you can respond accordingly.

Conclusion

The makeup of today's enterprises is changing, no doubt about it. Basic server/client infrastructures, the exclusive use of company-controlled assets and the physical containment of employees have long ceased to be the norms. Enterprise technology continues to grow in use and complexity. Workers are demanding more from their employers. All of this is leading to an ever-changing security perimeter. You have to ensure that your security teams are keeping pace and staying ahead of the attackers.

In this paper, we examined some of the challenges that organizations are currently facing. Our goal in evaluating today's challenges is not to cause angst or promote the fear that new technology is rendering security ineffective. Instead, we encourage your organization to leverage the changing landscape to find ways to make your security posture more adaptive—or, if you will, *more fluid*.

It is imperative that organizations today take every opportunity to consider the security implications of the complexities of their environment. For example, how many unknown devices are utilizing your networks; how much of your workforce, customer base and partner network is mobile; and how much of your enterprise exists in the cloud? By knowing answers to such questions, you begin to move to a state of viewing the intricacies of the environment as *opportunities* to create a stronger security posture rather than as hurdles.

Lastly, we also put forward some measures that you can use to further your security posture and environmental awareness. By taking advantage of continuous network and log monitoring to complement endpoint monitoring, security and operations teams can build holistic pictures of what their behavior should *normally* look like. Finding abnormal is much easier when you know what normal is. By building an adaptive, fluid security program, you are no longer reacting to incidents, but proactively and securely enabling the ever-changing organization.

About the Author

Matt Bromiley is a SANS Digital Forensics and Incident Response instructor, teaching Advanced Digital Forensics, Incident Response, and Threat Hunting (FOR508) and Advanced Network Forensics: Threat Hunting, Analysis, and Incident Response (FOR572), and a GIAC Advisory Board member. He is also a principal incident response consultant at a major incident response and forensic analysis company, combining experience in digital forensics, incident response/triage and log analytics. His skills include disk, database, memory and network forensics, as well as network security monitoring. Matt has worked with clients of all types and sizes, from multinational conglomerates to small, regional shops. He is passionate about learning, teaching and working on open source tools.

Sponsor

SANS would like to thank this paper's sponsor:

aruba
a Hewlett Packard Enterprise company

Last Updated: February 21st, 2019

Upcoming SANS Training

Click here to view a list of all SANS Courses

SANS Reno Tahoe 2019	Reno, NVUS	Feb 25, 2019 - Mar 02, 2019	Live Event
SANS Brussels February 2019	Brussels, BE	Feb 25, 2019 - Mar 02, 2019	Live Event
Open-Source Intelligence Summit & Training 2019	Alexandria, VAUS	Feb 25, 2019 - Mar 03, 2019	Live Event
SANS Baltimore Spring 2019	Baltimore, MDUS	Mar 02, 2019 - Mar 09, 2019	Live Event
SANS Training at RSA Conference 2019	San Francisco, CAUS	Mar 03, 2019 - Mar 04, 2019	Live Event
SANS Secure India 2019	Bangalore, IN	Mar 04, 2019 - Mar 09, 2019	Live Event
SANS St. Louis 2019	St. Louis, MOUS	Mar 11, 2019 - Mar 16, 2019	Live Event
SANS London March 2019	London, GB	Mar 11, 2019 - Mar 16, 2019	Live Event
SANS San Francisco Spring 2019	San Francisco, CAUS	Mar 11, 2019 - Mar 16, 2019	Live Event
SANS Secure Singapore 2019	Singapore, SG	Mar 11, 2019 - Mar 23, 2019	Live Event
ICS Security Summit & Training 2019	Orlando, FLUS	Mar 18, 2019 - Mar 25, 2019	Live Event
SANS SEC504 Paris March 2019 (in French)	Paris, FR	Mar 18, 2019 - Mar 23, 2019	Live Event
SANS Munich March 2019	Munich, DE	Mar 18, 2019 - Mar 23, 2019	Live Event
SANS Norfolk 2019	Norfolk, VAUS	Mar 18, 2019 - Mar 23, 2019	Live Event
SANS Secure Canberra 2019	Canberra, AU	Mar 18, 2019 - Mar 29, 2019	Live Event
SANS Doha March 2019	Doha, QA	Mar 23, 2019 - Mar 28, 2019	Live Event
SANS Jeddah March 2019	Jeddah, SA	Mar 23, 2019 - Mar 28, 2019	Live Event
SANS SEC560 Paris March 2019 (in French)	Paris, FR	Mar 25, 2019 - Mar 30, 2019	Live Event
SANS Madrid March 2019	Madrid, ES	Mar 25, 2019 - Mar 30, 2019	Live Event
SANS 2019	Orlando, FLUS	Apr 01, 2019 - Apr 08, 2019	Live Event
SANS Cyber Security Middle East Summit	Abu Dhabi, AE	Apr 04, 2019 - Apr 11, 2019	Live Event
SANS London April 2019	London, GB	Apr 08, 2019 - Apr 13, 2019	Live Event
Blue Team Summit & Training 2019	Louisville, KYUS	Apr 11, 2019 - Apr 18, 2019	Live Event
SANS Riyadh April 2019	Riyadh, SA	Apr 13, 2019 - Apr 18, 2019	Live Event
SANS Boston Spring 2019	Boston, MAUS	Apr 14, 2019 - Apr 19, 2019	Live Event
SANS Seattle Spring 2019	Seattle, WAUS	Apr 14, 2019 - Apr 19, 2019	Live Event
FOR498 Battlefield Forensics Beta 1	Arlington, VAUS	Apr 15, 2019 - Apr 20, 2019	Live Event
SANS FOR585 Madrid April 2019 (in Spanish)	Madrid, ES	Apr 22, 2019 - Apr 27, 2019	Live Event
SANS Northern Virginia- Alexandria 2019	Alexandria, VAUS	Apr 23, 2019 - Apr 28, 2019	Live Event
SANS Muscat April 2019	Muscat, OM	Apr 27, 2019 - May 02, 2019	Live Event
Cloud Security Summit & Training 2019	San Jose, CAUS	Apr 29, 2019 - May 06, 2019	Live Event
SANS Pen Test Austin 2019	Austin, TXUS	Apr 29, 2019 - May 04, 2019	Live Event
SANS Riyadh February 2019	OnlineSA	Feb 23, 2019 - Feb 28, 2019	Live Event
SANS OnDemand	Books & MP3s OnlyUS	Anytime	Self Paced