SANS

Analyst Program

Whitepaper

2022 SANS Protects: The Endpoint

Written by **Matt Bromiley**

April 2022

cisco SECURE

Introduction

Security teams are only as strong as their visibility and telemetry. Insight into an organization and its various assets is the primary enabler of effective detection and response capabilities. Endpoints, such as user workstations, laptops, servers, and cloud-based systems, are the most prevalent type of asset within organizations—and equally the ones that adversaries target most. Endpoints contain data, store user account credentials, and link to other parts of the network. A single endpoint can be an entry vector, a form of persistence, and an exfiltration point for adversaries.

As such, organizations work hard to protect their endpoints. One way that organizations help secure endpoints is through investments in endpoint-centric technologies, including traditional antivirus and endpoint detection and response (EDR) tools. Security postures typically center around endpoint technology *first* and often include other types of telemetry to *support* endpoints. However, despite the investments that many organizations have made, adversaries still perform successful intrusions.

It's time to consider what capabilities our endpoint defenses have and whether security teams are utilizing them to the fullest potential. In this SANS Protects paper, we look at threats to endpoints and ways that organizations can overcome or mitigate them. Our SANS Protects papers focus on threats and mitigations, helping organizations consider elements of security that they should be implementing.

Deploying an endpoint solution is no easy feat. Security teams have to consider many factors, including system resource utilization, how to consume and act on data, and ways an endpoint solution will impact users. Regardless, organizations must choose a solution that will best enable their security team to deliver on the requirement of protecting the organization, its data, and its users. Some key considerations for endpoint security, especially on the heels of today's threats, include:

- Endpoint solutions should take advantage of newer technology trends, such as AI/ML, advanced detections, and moving target defenses (MTD).
- Endpoint solutions should inform and enable security teams. Simply reporting detections, without additional context or associated telemetry, does not provide teams the necessary advantage.
- Endpoint security should be coupled with detection and response capabilities, so that teams can triage, analyze, contain, and block with the same platform.

Our SANS Protects series is also meant to be thought-provoking. As you work your way through this paper, we encourage you to evaluate the current state of endpoint security within your organization. It is possible—we hope *likely*—that you already have an endpoint security solution deployed. In that case, we encourage you to explore what has been deployed within your organization and confirm that you are receiving the protections your security team is making assumptions on.

Threats to the Endpoint

Endpoints represent a unique threat to organizations because they are:

- The most common asset type within an enterprise
- The most sought-after target by adversaries for various stages of an attack
- Possibly the *target* of multiple types of threats that other security controls may be expected to *prevent*

For example, consider a spear phishing email opened on a user's desktop. While we would expect email security controls to prevent delivery of the spear phishing email, it is now a threat that the endpoint must detect and block. Unfortunately, some adversaries rely on this complexity, crafting emails to evade both email and endpoint defenses.

This does not give an excuse for inadequate email security; rather, it expands the threat profile for endpoints within an organization. We believe this gives reason for a strong EDR capability. Without it, security teams cannot expect to stand up against adversaries. Security teams are the final line of defense, from a technology perspective. With this in mind, we will look at some of the most popular threats to today's endpoints.

Final Line of Defense

One of the largest threats to endpoints stems from the fact that endpoints are, as briefly mentioned, the final line of defense. Spear phishing emails are meant to execute on an endpoint. Malicious documents, spreadsheets, PDFs, and advertisements are all intended to be opened or viewed by a user on an endpoint. However, these threats often pass through a separate line of security controls before arriving at the endpoint. Unfortunately, if they are being addressed at the endpoint, then other controls have likely failed to block or intercept them.

While this may not be a threat in the traditional sense, this final line of defense can create complexities for endpoint security products. EDR tools must be able to receive and analyze telemetry related to network traffic, emails, process execution, file permissions, system changes, user behavior, and much more. The more complex an endpoint agent is, the more complex detections that security teams can write. However, this may also drain system resources and/or slow the user experience.

Furthermore, endpoint defense can establish a false sense of protection for security teams. Relying on the idea that "the endpoint agent will catch it" is not a valid assessment of other security controls. "Why did it get as far as the endpoint?" may be a better question and one that can help assess the value of pre-endpoint controls that may be in place within the environment.

Vulnerabilities

Another threat that has plagued many organizations over the years has been vulnerabilities. Given the function of endpoints, this threat category can be tough to classify and protect. Vulnerabilities may pertain to the operating system, a piece of hardware within the physical system, or a third-party piece of software. As recent times have shown, sometimes the vulnerabilities are within *libraries* built into the third-party software, which then put the endpoint at risk.

The function of an endpoint can also modify its risk profile. For example, imagine an external-facing server versus an internal, air-gapped system. From an endpoint perspective, they may be the same type of asset but their risk profile is very different from a vulnerability perspective. Adversaries may look to find holes in the external-facing software, such as a web server, supporting library, or open port. As seen in recent vulnerability disclosures, adversaries waste absolutely no time in mass scanning the internet when a new vulnerability is released.

The other risk vulnerabilities pose is their capability to increase adversaries' success rates. For example, consider an adversary who has gained a small footprint inside an organization but is unable to move further due to a strong defense-in-depth strategy. Perhaps lateral movement capabilities are locked down or accounts are granted least permissions—until adversaries discover an internal vulnerable system that allows them to pivot or steal credentials. Now, a good strategy of network segmentation has been thwarted by an internal-only, however outdated, system.

Post-Exploitation Toolkits

One threat that often impacts endpoints is the use of widely known and well-reputed post-exploitation kits. Of course, the first that come to mind are the ever-present Cobalt Strike and Metasploit. While marketed as adversary simulation or offensive security tools, all-in-one post-exploitation kits are some of the most pervasive adversary tools. A June 2021 article from Threatpost reported that Cobalt Strike usage was up 161% year-over-year in cyberattacks,[1] with the tool used by both APT and general-commodity or ransomware actors.

Despite being created for legitimate or research uses only, proof-of-concept exploits and post-exploitation kits represent a tricky double-edged sword for security teams. While they are used by legitimate red teamers, adversaries utilize the exact same tools to launch attacks against organizations. Given the scope and usage of these tools, it is prudent to incorporate them into any endpoint threat model.

Post-exploitation toolkits also represent a significant threat to an organization's endpoints because they come bundled with dozens (if not hundreds) of exploitation capabilities. Adversaries can gain initial access on a system and rotate through dozens of attacks, exploits, privilege escalation, and lateral movement attempts until they find something that works.

These toolkits also come with an element of automation, something we've seen multiple threat actors use to their benefit. Cobalt Strike, for example, has the capability for adversaries to load *Aggressor Scripts*, which are automated series of actions that occur on a victim system. This can give adversaries a significant advantage if they can assemble a script that takes them from 0 to 100 in a minimal amount of time.

[1] "Cobalt Strike Usage Explodes Among Cybercrooks," Threatpost, https://threatpost.com/cobalt-strike-cybercrooks/167368/

Unnecessary User Entitlements

Another high-priority threat to the endpoint, and one that adversaries always seek to take advantage of, are the inherent permissions that follow user accounts as users go about their normal day. Such permissions include logging into websites or checking email. For example, all the actions that occur on a Windows system are tied to an account. Whether it's running a web browser, opening a PDF, or having a video conference, user permissions govern what users can and cannot do.

Additionally, as normal domain operations take place, accounts move in and out of systems without any interaction from the user. Servers may go without an interactive session for weeks or months, yet accounts log in and out of systems hourly or daily. Regrettably, users are often given far too many permissions.

Perhaps one of the most formidable threats to organizations is ransomware. Targeted attacks by skilled and resourced adversaries can result in locking up an organization, costing millions of dollars and resulting in public embarrassment and halted operations. However, organizations should remember that ransomware is a combination of techniques exploited by adversaries. When mitigated, organizations can put a serious dent in adversaries' success rates.

When accounts—whether local, domain user, or service—are provisioned, there is often a specific set of permissions that accounts need be given. However, because user access policies can be difficult to manage and maintain, organizations give users far more permissions than necessary. For example, rather than limit a user in the accounting department to the accounting-specific files, the user may be given access to the entire file share. Adversaries love these situations and quickly move to take advantage of them.

User permissions are often easy to steal from infected workstations. Whether it's finding or guessing the account password, escalating via an exploit, or impersonating an account, post-exploitation toolkits often include automated or single-click means to gain user credentials. Adversaries can quickly deploy them and go from a single infected process to an infected system—or worse, an infected domain.

In-Memory Exploitation/Fileless Malware

If this paper were written several years ago, we might have focused our discussion of endpoint security on topics such as next-gen antivirus, on-disk malware detection, or the use of indicators of compromise (IoCs) to identify static, malicious files. However, adversaries, offensive security tool creators, and malware authors have all realized that disk-based detections are reliable enough that robust versions can identify various permutations. Furthermore, many endpoint solutions are based on scanning and analyzing on-disk files, leaving a system's memory a wide-open playground.

To take advantage of this, it is now common to see the usage of in-memory, or fileless, malware. Rather than leave static artifacts in disk, in-memory malware exists only in memory space to evade traditional disk-based detection techniques. This can also complicate post-incident forensic analysis, which may rely on more traditional disk artifacts to triage incidents.

In-memory, or fileless, malware can complicate detections because memory space on a victim system can be sizeable and/or contain a significant amount of data. Luckily, despite leaving less disk-based evidence, adversaries often go after the same processes or items in memory (such as account tokens), and thus we can wrap protections around common items within memory, rather than try to emulate and monitor the entire memory space.

Endpoint Agent Evasion and Tampering

Finally, another significant threat facing endpoints is meddling with the endpoint itself. While the concept of endpoint evasion may refer to a technique that avoids detection mechanisms, some adversaries seek to tamper with or disable the endpoints entirely. This technique presents a challenge to security teams because they may lose all insight into an infected system(s). Detection and response programs built around an agent's visibility may be stuck without a solution if they lose the agent entirely.

However, this process is often easier said than done. Many endpoint products come with anti-tampering capabilities.

Protecting the Endpoint

While the preceding discussion of vulnerabilities does not encompass every single threat facing an endpoint, it summarizes key threats that adversaries are utilizing today to gain a foothold, escalate privileges, and move around a compromised environment. Because the threat profile for endpoints is so large, organizations should look to robust endpoint defenses that (1) provide visibility and telemetry, (2) allow for detection of multiple types of tactics, techniques, and procedures (TTPs), and (3) include response capabilities.

We assume that any endpoint defense strategy stems from an EDR platform, which typically includes many of the advanced features that we expect organizations to harness. Simple antivirus solutions, while useful in detecting and preventing low-level threats, cannot provide the type of advanced analytics and forensic response capabilities that security teams need.

Asset Visibility and Recognition

Any good security program requires asset visibility and recognition. Endpoints are no different. In fact, endpoint security and asset visibility are inextricably linked. The endpoint agent that provides additional detections also should first provide visibility into said asset. This is an easy way for security teams to ensure that they have full endpoint deployment and that assets are reporting in as expected. An anomaly in asset visibility is already a security concern and may be an early indicator of suspicious activity on a system.

When we talk about asset recognition, we ask that endpoint monitoring go a step above simply checking in to a central console. Endpoint visibility should also include key system statistics, such as:

- Running processes, outside of detections
- Installed software
- Network details connections
- Logged-in users
- Asset metadata, such as grouping, geolocations, policy details, etc.

An asset inventory list coupled with the metadata points already place a security team in a well-balanced position to keep an eye on their environment. Furthermore, these metadata can also be used to write detections themselves, outside of the other built-in functionality we'd expect from an EDR tool.

Automated Prevention, Response, and Blocking Capabilities

When it comes to endpoint defenses, detection and response should be the beginning—or an assumed, baseline capability. However, we'd like to see endpoint defenses be more advanced than just detection or reporting. Endpoint defenses should also come equipped with automated capabilities, enabling security teams to create and deploy playbooks in response to certain TTPs. This allows security teams to build high-fidelity *reactions* to adversary actions.

We realize that some automated capabilities are often handled at the platform level, not the endpoint specifically. This is a fair trade-off, and perhaps a better one because it allows for centralized management, creation, and deployment. Furthermore, playbooks can be specified for certain types of endpoints or groups, providing highly granular control (a welcome thing in security). However, organizations should inquire about the latency and requirements for a platform to automatically respond *on behalf* of an endpoint, based on detected activity. The worst situation security teams can be in is having a false sense of security, thinking that they deployed protections or put preventions in place only to find out that they are not working as expected. No one likes a false sense of security!

Endpoint++: An XDR Strategy

The best forward-thinking strategy is a recognition that while endpoints are a category unto their own and deserve their own specific/custom protections, no environment consists of strictly endpoints. For this reason, incorporating endpoint defenses with other telemetry, such as network or email data, provides a robust, multisource capability that may be what security teams need. This strategy, known as *extended detection and response (XDR)*, brings together multiple sources and really empowers security teams to think about the impact that threats have on their environment.

Adversaries and their associated TTPs seldom touch a single source of telemetry. Lateral movement requires activity on the *network* and the *endpoint*. Deployment of ransomware impacts file shares and users, all connected via the network. Malware delivered via an email lands on a system and begins beaconing out. The realization that artifacts crisscross all the time helps security teams combine and correlate multiple sources of telemetry for a truly robust anti-adversary security posture. While we may not want to collect network telemetry *from* an endpoint, we can easily collect it from our environment and *combine* it for easy analysis.

Furthermore, an XDR strategy brings in the automated actions already discussed.

Case Study: Fileless, In-Memory Ransomware Attack

Our case study looks at one of today's biggest threats to organizations—and thus endpoints: ransomware. As mentioned earlier, ransomware is a unique situation for security teams because it focuses on the *objective* of the adversary as the goal, which typically is *after* a collection of TTPs. Let's examine this further.

Ransomware attacks, as shown in Figure 1, often begin with an entry vector. This may consist of a spear phishing email, drive-by download, or exploited vulnerability. Regardless of the delivery mechanism, it is now a matter of minutes before a malicious payload moves from deliver to in-memory operations. One collection of TTPs used a lot by adversaries and post-exploit toolkits these days are fileless, or in-memory, attacks, which leave little evidence on disk.

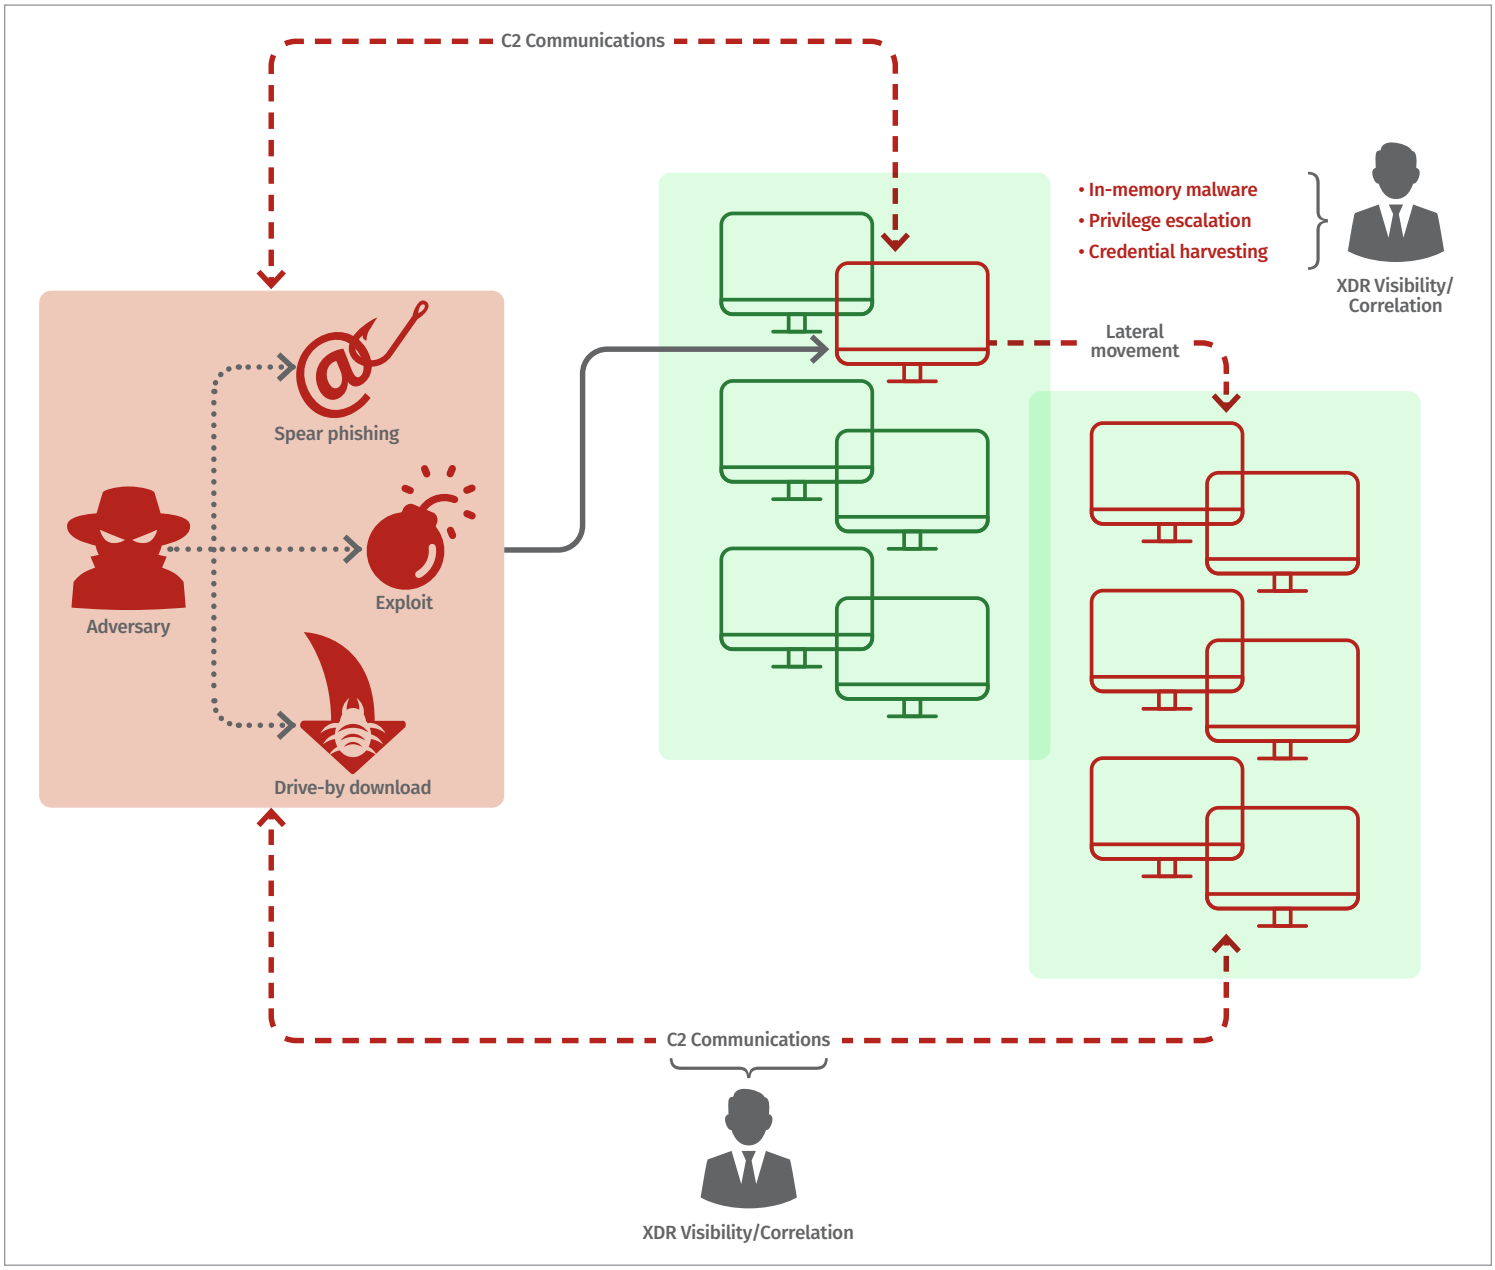

Figure 1. Ransomware Attack

Legacy endpoint protection tools will miss modern attacks. They focus on on-disk artifacts and look for static indicators, such as a bad hash, filename, location, or a combination of the three. In-memory malware avoids disk-based artifacts. Thus, we look to advanced endpoint detection capabilities, which can monitor in-memory operations, processes, system calls, network connections, and more. A capable endpoint agent can provide a wealth of telemetry back to security teams so that they can effectively triage when an incident is detected.

Let's now assume that our adversary has found some success in their attack, as shown in Figure 1. A detection-only or limited-capability endpoint might stop at this point. Providing alerts with minimal context does little for security teams. Instead, we'd rather see an endpoint agent not only provide robust telemetry, but also have inherent response capabilities. Can we use the endpoint agent to block processes, shut off network connections, or quarantine the entire system?

Going a step further, we need to triage this incident while monitoring for additional bad activity within the environment. An EDR or XDR platform should be able to take the TTPs observed during the attack and quickly pivot to look for other malicious activity. This level of automation allows organizations to scale an attack on *one* system to a detection for *many* systems, as well as limit the adversary's ability to go elsewhere or cause additional damage.

Conclusion

Endpoints represent one of the largest attack surfaces for organizations. Ranging from user workstations and laptops to servers, either cloud or on-prem, endpoints are the key "systems" within any enterprise. This places endpoints at the top of any adversary's target list, as they seek to enter and maintain presence in a victim organization. But this priority works both ways. Endpoint security has reached a point where organizations can build robust security postures that rely on endpoint visibility, detection, and response capabilities.

In this SANS Protects paper, we explored threats to and protections available for enterprise endpoints. With endpoints being a primary target for adversaries and the place where users are the most, endpoint defenses are a critical part of any security approach and a necessary component for security teams to effectively handle incidents within their environment. Some of our key takeaways include:

- Look for multipurpose endpoint platforms that assist organizations with detection and response to incidents.
- Prepare robust, granular detection capabilities to ensure that even wily adversaries will have a tough time defeating defense.
- Let endpoint management double-up to increase visibility and catalog endpoints.
- Utilize advanced technology and analytic capabilities to detect anomalies and evasion attempts.

Regardless of what technology may be implemented, adversaries will still seek ways to evade defenses and gain footholds in a victim environment. Thus, while endpoint defenses represent a large footprint, they should be a portion of security monitoring, not the entire plan. As stressed earlier, adversaries who disable or tamper with endpoint defenses should not be able to eliminate visibility. Instead, endpoints should be part of a robust, multisource approach that utilizes the advanced features of endpoint defenses, coupled with additional data points, to quickly detect adversaries in an environment.

Sponsor

SANS would like to thank this paper's sponsor: