

Whitepaper

Expanding the Security Toolbox

Written by **Matt Bromiley** June 2021



©2021 SANS™ Institute

From a cybersecurity perspective, defending an organization can seem like a daunting, and sometimes impossible, task. As businesses embrace more technologies, threat actors waste little time discovering new tactics and techniques to infiltrate organizations. Smarter and more evasive than ever, threat actors have shown breach after breach that they are not afraid to use multiple tools to achieve their objectives.

Their persistence can seem like an insurmountable problem to defenders, many of whom struggle to keep up with the ever-changing threat landscape. Note this important lesson, however, for today's attacks: Just as threat actors can switch tactics when needed, defenders also have a vast toolbox and access to data. The problem is that many seldom use what is available to them or know the multiple ways in which to use that data.

Many in information security are aware of the data available to them for analysis but fail to capitalize on what is within their reach. We wondered: What is the roadblock? Why, for example, do some organizations rely heavily on endpoint-centric defenses and ignore the plethora of network data they could integrate? For many, the issue comes down to recognizing, combining and correlating valuable data—no easy feat. Unfortunately, many ignore this task while the problem grows.

In this whitepaper, we tackle this problem head-on. Today's technologies make it much easier to combine and correlate disparate data sources.

Furthermore, as examined here, when multiple data sources come together, an organization can write better-contextualized, less-brittle detections and detect, possibly even stop, threat actors earlier in a breach. Multisource detections also provide excellent audit trails for incident response, allowing you to conduct and conclude investigations faster. With an expanded toolbox, organizations can take on previously insurmountable security challenges.

Access to this expanded toolbox allows the security team to rely on threat libraries—such as the MITRE ATT&CK® Matrix—to evaluate how threat actors utilize techniques and to then implement effective defenses. Rather than continually chasing alerts and hoping yesterday's detections will catch tomorrow's threat actors, security teams now have the tools and frameworks at their disposal to detect threats earlier and respond appropriately.

This paper focuses on the power of combining multiple data points to achieve as much visibility as possible within an enterprise. When coupled with an understanding of threat techniques, a security team will see their detection and response capabilities grow exponentially. As you read this paper, consider the following questions about your own environment:

- How much visibility do we have into the various elements of the organization?
- What data points does my security team currently utilize to detect and respond to incidents?
- Does my security team write their own detections? If so, do we utilize all the data points identified above?
- When we consider our risk exposure to certain attacks, do we assess how we can increase data visibility to decrease risk?

These questions, and many more that come up in this paper, enable you to evaluate the current state of your security team and determine what changes you can make to ensure your team takes full advantage of all the data points available to it.

Why do some organizations rely heavily on endpoint-centric defenses and ignore the plethora of network data they could integrate?

When Our Powers Combine

The concept of expanding the security toolbox begins with one word: *visibility*. Visibility is critical for effective enterprise defense and is much easier said than done, especially given the vast reach of many organizations today. A modern organization could easily have operations both on-premises and in the cloud, taking advantage

of architectures such as serverless and containerization across multiple operating systems in both. Simultaneously, the organization may be spread across multiple geographic areas and with employees working remotely, further complicating an architecture that relies on central access points. Such conditions can create gaps in visibility that result in undetected issues.

The security team should strive to gain as much visibility as possible over its assets, which it can do in multiple ways: from collecting data directly from a source to utilizing a third-party technology, such as endpoint detection and response (EDR), to generate and forward security-centric events. The more security-centric a data source is, the faster the security team can derive value from it. Figure 1 looks at the most common technologies within the modern enterprise and types of visibility available for the security team.

The various sources of telemetry in Figure 1 represent only a subset of the various data types out there, but even this

subset confirms one thing: Your security team has multiple options for gaining insight into various elements of the organization. Regardless of the size of each technology with respect to the overall environment, you want to leave no assets unseen.

Visibility, however, encompasses much more than simply "seeing" a particular asset or asset class. Visibility also facilitates identification of the data sources the security team can use to detect and respond to incidents within the environment—along with *how easily an attacker can evade a particular source of telemetry*. Security teams that base both their visibility *and* detections on a single source of telemetry risk creating an exponential single point of failure. Threat actors need only evade a single security control, allowing them to achieve their objectives much more easily. Furthermore, without visibility, we cannot associate multiple evasive behaviors with each other

Source	Type of Telemetry
Endpoints	System logs
	• EDR
Networks	Packet capture (PCAP)
	NetFlow
	• Network detection and response (NDR)
	Protocol logs
	Network device logs
Cloud-hosted assets/services	• Asset logs
	Provider network data
	• Similar telemetry from above sources
Third-party applications	• Vendor logs
	Network traffic
Containers	 Container/application logs
	• Filtered EDR
	Network logs
Threat intelligence	Input feeds
	• Internal incident data
User/authentication logs	 User logins/logouts
	Authenticated events
	Account usage

Figure 1. Common Technologies and Associated Visibility

Wily attackers can thwart even the best-laid security plans. Visibility is crucial to effective defense—not only because it encapsulates your entire environment but also because it guarantees an attacker cannot go completely "dark." Once we have the data in one place, how do we write detections that take advantage of multiple telemetry sources?
 The first question is easily answered by advancements in technology and correlative data platforms that allow for security teams to easily ingest multiple sources of data. Your organization could easily generate gigabytes of logs per hour. Even homegrown, open source solutions could handle these volumes.

• How do we bring all these disparate data sources together?

The second question requires us to utilize third-party sources and threat intelligence to gain insight into threat actor behavior. Efficient detections begin with understanding how a threat actor employs a certain technique and aligning that knowledge with your various telemetry points. It is no secret that the best resource to bring these concepts together is the ATT&CK Matrix.

Aligning Detections with MITRE ATT&CK

respond with the same two questions:

As previously mentioned, multiple sources of telemetry raise a team's confidence that even if an attacker were to evade one detection, a "back-up" source still exists to catch them in the act. This confidence also allows a security team to rely on multisource detections, meaning they do not spend every day chasing a single source of telemetry to keep the environment secure. Additionally, multisource detections allow for better coverage of the multitude of threat actors in the wild.

You must update your security team's visibility to include all the technologies you utilize. Attackers redefine the threat landscape daily, making an asset that was safe today vulnerable tomorrow. Multiple points of visibility offer resiliency that attackers cannot easily evade.

For example, consider an environment that has a security program almost entirely rooted in endpoint telemetry. Whether the team ingests system logs or utilizes security-centric EDR data, the outcome remains the same: reliance on a single source of data. Visibility is determined by whether an agent is "checking in" or the system is shipping logs, neither of which confirms that the system is powered on and uninfected. If an attacker were to take advantage of such a system, what mechanisms does the organization have in place to alert the security team that something has gone wrong? The team would get stuck in an all-too-familiar loop, checking agents daily to confirm that systems are still visible and *reacting* to incidents rather than crafting detections to get ahead of them.

Instead, consider an environment that utilizes both endpoint and network telemetry. An endpoint agent may stop checking in, but the security team maintains some visibility via network traffic. If an attacker were to execute malicious code or connect outbound, the security team would still have network visibility despite a failed endpoint agent. Similarly, a cloud-hosted asset may not permit an endpoint agent to be installed, but a security team can combine back-end operational logs with exported network flow data and achieve the same type of resiliency.

What we are suggesting is not a new concept. Information security professionals have been pushing for enterprisewide visibility for many years now. Organizations often

Analyst Program I Expanding the Security Toolbox

BITS Jobs

Adversaries may abuse BITS jobs to persistently execute or clean up after malicious payloads. Windows Background Intelligent Transfer Service (BITS) is a low-bandwidth, asynchronous file transfer mechanism exposed through Component Object Model (COM).^{[1][2]} BITS is commonly used by updaters, messengers, and other applications preferred to operate in the background (using available idle bandwidth) without interrupting other networked applications. File transfer tasks are implemented as BITS jobs, which contain a queue of one or more file operations.

The interface to create and manage BITS jobs is accessible through PowerShell and the BITSAdmin tool.^{[2][3]}

Adversaries may abuse BITS to download, execute, and even clean up after running malicious code. BITS tasks are self-contained in the BITS job database, without new files or registry modifications, and often permitted by host firewalls.^{[4][5][6]} BITS enabled execution may also enable persistence by creating long-standing jobs (the default maximum lifetime is 90 days and extendable) or invoking an arbitrary program when a job completes or errors (including after system reboots).^{[7][4]}

BITS upload functionalities can also be used to perform Exfiltration Over Alternative Protocol.^[4]

ID: T1197

Sub-techniques: No sub-techniques

- (i) Tactics: Defense Evasion, Persistence
- (i) Platforms: Windows
- ① Permissions Required: Administrator, SYSTEM, User
- ① Data Sources: Command: Command Execution, Network Traffic: Network Connection Creation, Process: Process Creation, Service: Service Metadata
- ① Defense Bypassed: Firewall, Host forensic analysis

Contributors: Brent Murphy, Elastic; David French, Elastic; Red Canary; Ricardo Dias Version: 1.2

Created: 18 April 2018

Last Modified: 13 April 2021

To help illustrate the value of a multisource detection, let's examine a single technique, BITS jobs from the ATT&CK Matrix, to prove this concept. Figure 2 provides a snippet of the Background Intelligent Transfer Service (BITS) jobs technique detail, ID T1197, from ATT&CK¹

As shown in Figure 2, a BITS job is a technique that allows an attacker to abuse a native system tool (**bitsadmin.exe**) to download and upload data from an infected system. Multiple threat groups use this technique, including at least APT39, APT41, Leviathan and Patchwork. We can also glean from ATT&CK that downloading via BITSAdmin is integrated into post-exploitation frameworks such as Cobalt Strike. Immediately, the security team should realize this: An attacker's abuse of BITS jobs is not a standalone or single-actor technique and, therefore, awareness of such can expose multiple attackers and popular exploit kits.

Let's craft a detection. A security team would likely default to crafting an endpoint-centric detection to identify the abuse of an on-disk executable included with the Windows operating system. That is a solid place to start, and MITRE provides us insight that most BITS jobs abuse comes from the invocation of **bitsadmin.exe** on a Windows system. This knowledge allows the security

team to consider various endpoint telemetry sources as its primary source of detection. The team could easily double up on multiple sources of endpoint telemetry, extracting EDR data *and* processing execution system logs. Figure 2. Screenshot from the ATT&CK Matrix for T1197, BITS Jobs

When crafting detections, consider using as many sources of telemetry available to your security team as possible. Multisource detections allow you to increase the likelihood of true positives and provide wider coverage of threat groups and malicious tools.

¹ BITS Jobs, Technique T1197 - Enterprise | MITRE ATT&CK, https://attack.mitre.org/techniques/T1197

The next step involves crafting a detection by examining how a threat actor utilizes a technique. Using ATT&CK as our library and guide, we can see how one threat actor (named Patchwork)



utilizes BITS jobs to download malicious files. Knowledge of command invocation proves instrumental in crafting an effective endpoint detection. See Figure 3.

Security teams with a single source of telemetry would stop here. Without being able to extend their detections into other data points, they could easily detect malicious BITS

usage—assuming the attacker does not disable defenses first. However, the BITS tool utilizes network connectivity. It inherently straddles both endpoint and network detections. Teams with access to both sources of telemetry would find that in BITSAdmin HTTP requests there exists another element for detection: the Microsoft BITS/7.5 user-agent. See Figure 4.

Request HEAD /m/mb.exe HTTP/1.1 Connection: Keep-Alive Accept: */* Accept-Encoding: identity User-Agent: Microsoft BITS/7.5 Host: 92.63.197.48

Attackers could potentially evade an endpoint defense, but they would be hard pressed to evade the network. Armed with this knowledge, a resourceful security team could craft a multisource detection that looks at network and/or endpoint telemetry and provides a higher fidelity into malicious activity. Figure 5 shows an example of such a detection Figure 4. Snippet of a Malware Report Showing BITSAdmin Network Connectivity HTTP Headers³

Report on the "Patchwork" Threat

Actor and Their Usage of BITS Jobs²

using Elastic's Event Query Language (EQL).

Note that in Figure 5, our ideal state, a security team can detect across multiple datasets simultaneously. Process execution and network connectivity are two of many artifacts that can tell the same story but from sequence by host.name with maxspan=30s
[process where process.name : "bitsadmin.exe" and event.type == "start"]
[network where user_agent.original : "Microsoft BITS*" and not
cidrmatch(destination.ip, "10.0.0.0/8", "172.16.0.0/12", "192.168.0.0/16",
"127.0.0.0/8")]

a different viewpoint. By crafting this single, two-part detection with an **OR** statement, we can allow for evasion techniques that an attacker may employ. Furthermore, the network detection allows for additional flexibility: The security team can choose to ignore internal

Figure 5. Snippet of an EQL Query That Looks Across Multiple Evidence Sources to Identify Suspicious Activity

IP addresses, limiting potential false positives and increasing the fidelity of the detection.

With more data sources comes more opportunities for false positives. As such, make sure your multisource detections account for network directionality, internal vs. external subnets, normal system execution activity and expected behavior within your organization. Any chance to tune and limit false positives will provide higher fidelity and give your security team a faster reaction time.

² Unit42, "Updated BackConfig Malware Targeting Government and Military Organizations in South Asia,"

https://unit42.paloaltonetworks.com/updated-backconfig-malware-targeting-government-and-military-organizations ³ Hybrid Analysis, "IMG033435056_2018-JPG.js,"

www.hybrid-analysis.com/sample/e79536d1d95ebe24eea16e18af581dd6918f8af67e380ee5a7949ca2894db316/5c329fd67ca3e14811376756 (Free automated malware analysis service powered by Falcon Sandbox. View of online file analysis results.)

Enhancing Incident Response Operations

Multisource detections are not the only benefit of our proposed approach. By combining multiple data sources to build *better detections*, security teams will also find themselves with better data for incident response. In fact, bringing multiple sources of data together to identify attacker activity is a classic incident response (IR) undertaking. We simply want to perform this earlier and utilize it to build detections.

If an incident does occur, a security team with multiple correlated data points will enjoy shorter investigations, easier remediation and quicker return to business as normal. Throughout the two case studies we examine in the next section, consider whether you have experienced any of the techniques we call out. If so, how long did it take your organization to detect, respond and recover from these incidents? Imagine if we had multiple sources not only for detection but also to quickly scope and contain active incidents.

Case Studies

Joining multiple data sources for incident detection and response is often easier said than done. Even in environments where data sources peacefully coexist, however, defenders may lack the skills or knowledge to combine and analyze correlated datasets. In the following case studies, we examine a series of threat actor techniques and the data sources that you can correlate for effective detection and response.

Case Study 1: Credential Harvesting and Lateral Movement

There is likely no more common trio of tactics and techniques than credential harvesting, system discovery and lateral movement. Threat actors exploit this trio in attacks ranging from long-term advanced persistent threat (APT) intrusions to short-lived ransomware heists. After all, threat actors hardly ever conduct single-system attacks and land on the exact system with the perfect compromises.

Instead, threat actors understand that they must pace themselves through an environment. This process usually begins with gaining an initial foothold, then escalating privileges, and after stealing credentials they can start moving laterally throughout the environment. In between each of those actions, they may execute multiple system discovery and/or reconnaissance commands to identify additional accessible devices within the environment.

The more actions a threat actor performs, the more multisource detections we must implement. Let's examine a simple depiction of this scenario, provided in Figure 6 on the next page.



Multiple chances exist for endpoint detections, but threat actors can easily thwart them. Instead, each step of the attack provides multiple points of detection:

- **1.** An infected system, used to steal credentials *and* look for additional systems in the environment, likely casts out lots of traffic and generates lots of telemetry. You can use each source for either incident detection or response.
- 2. If an attacker moves to an additional system, network telemetry will capture the movement, while credential usage and endpoint telemetry will capture an authentication event. You can tie together all this information and correlate it to detect lateral movement by a likely privileged account.
- **3.** Attackers looking to steal data or make an impact will seldom stop at user workstations. Therefore, the discovery/reconnaissance process begins again. An attacker looking for servers, using a particular account, creates the same types of noise in account usage and network-based searching.
- **4.** When an attacker moves to a system within a different part of the network, multiple points of telemetry once again provide backup methods for detecting malicious activity.

The above describes a simplified example, but it provides a solid first start to understand the power of multisource detections. Walking back through the preceding example, even if you remove one (or in some cases, two!) sources of telemetry, another takes its place. Removing an endpoint agent from the systems, for example, still leaves suspicious network traffic and account usage. Figure 6. High-Level, Simplified Depiction of Common Techniques Used by Attackers Within Enterprise Environments

Case Study 2: Encrypted and Exfiltrated Data as Part of a Ransom Attack

If ever a piece or malware should bring worry into the minds of security analysts, ransomware is a top contender. In the first half of 2021 alone, we have seen ransomware attacks bring down healthcare, education and critical infrastructure networks. Ransomware attacks have collected hundreds of millions of dollars in ransom payments, meaning their strategy is working and that these threats are likely to persist for quite some time.

Ransomware attackers have recently upped the ante by exfiltrating data and using it to bargain for payment, meaning analysts must work to detect and stop these attacks more than ever. However, analysts should not think of ransomware as a standalone or single threat. Instead, ransomware refers to a collection of techniques that ultimately *leads* to data encryption and exfiltration. In fact, data encryption and exfiltration represent the *last* stages you want to alert on. Instead, we can look to multisource detections to find an attack much earlier in the attack life cycle. Figure 7 provides a walkthrough of this type of attack.



Figure 7. High-Level, Simplified Depiction of Common Techniques Used by Attackers Within Enterprise Environments Again, we see opportunities for multiple points of detection:

- Ransomware attacks often begin from an entry vector such as spear phishing or weak remote access. First-stage droppers may reach out and download additional files from a compromised IP address (Stage 1 C2). Malicious code and external connectivity provide detection opportunities from an endpoint and network perspective.
- 2. After the first stage of compromise, the attacker may follow up with credential harvesting and internal reconnaissance. We covered detections for these in our previous case study. Repeated techniques like this are excellent opportunities to deploy multisource detections that can catch myriad attackers and malware.
- **3.** After deploying initial malware and gaining a foothold in the environment, attackers will often include or transition their attack to a more robust C2 infrastructure. This action may be in anticipation of the upcoming ransomware stages. The transition provides multisource network evidence, ranging from HTTP and DNS logs to NetFlow and/or PCAPs.
- **4.** As previously mentioned, lateral movement and account usage trigger multiple points of telemetry for detections.
- **5.** When a ransomware attack identifies systems of interest, threat actors prepare the next stage of the attack to lock up and exfiltrate data. A significant amount of concurrent activity makes it difficult for attackers to remain truly silent. Between C2 network beacons (step 6), data uploads, mass file encryption and user account lateral movement, we can bring together almost every source the security team has access to for robust detection of these events.

You may notice that the title of this case study referred to data encryption and exfiltration. However, we diagrammed much more of the attack. Analysts should be asking this question when they prepare detections: At what stage of the attack life cycle are we finding the attacker? Data encryption and exfiltration is *late* in a ransomware attack. At that point, attackers have already had privileged access to your environment and have nearly completed their objectives. This is not the time to build detections (although, in the absence of anything else, they are a good starting point).

Instead, by looking at the attack, we could find ways to detect malicious activity earlier in the attack. Furthermore, we quickly found that our previous case study on credential harvesting and lateral movement was also applicable to our ransomware attack analysis. This discovery was not coincidental: Good detections that look for common attacker techniques can perform sweeping damage to attacker success because techniques are so often recycled among threat actors.

Conclusion

Too often, it can feel like attackers have the advantage when it comes to cybersecurity. Defenders constantly play catch-up to attacker techniques as more and more breaches become public. As soon as a security team can get its hands on tooling to assist them in their goals, attackers come up with multiple ways to evade and subvert specific technologies. Organizations that rely on a single point of telemetry for detections have instituted a single point of failure—and attackers are waiting for chances to evade detection. It can feel like a losing battle, but it really isn't. It is time for a change, and we argue that organizations can make an impact sooner rather than later.

Organizations that rely on a single point of telemetry for detections have instituted a single point of failure—and attackers are waiting for chances to evade detection.

In this whitepaper, we discussed the need for organizations to incorporate multisource detections in their toolkits. Attacker techniques, no matter how evasive, are never completely silent—and you may already have the data to detect them.

We also examined multiple examples and case studies that identified areas of success for multisource detections. Whether it is a single technique, such as utilizing BITS jobs to download a piece of malware or a complex, multistep attack like ransomware, myriad opportunities for success exist by detecting with multiple sources. Network traffic, endpoint telemetry, authentication logs, cloud provider audit trails and external threat intelligence are just a few examples that, when brought together, create an almost impenetrable trail of evidence that attackers will be unable to evade. Furthermore, you likely already have a lot of this data, and so you merely have the task of bringing it all together. The technology to combine and correlate these data points is now easily available and waiting to help your team to make an impact.

About the Author

Matt Bromiley is a SANS digital forensics and incident response instructor, teaching FOR508: Advanced Incident Response, Threat Hunting, and Digital Forensics and FOR572: Advanced Network Forensics: Threat Hunting, Analysis, and Incident Response. He is also an IR consultant at a global IR and forensic analysis company, combining experience in digital forensics, log analytics, and incident response and management. His skills include disk, database, memory and network forensics; incident management; threat intelligence and network security monitoring. Matt has worked with organizations of all shapes 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:



