cybereason | WHITEPAPER

# Operation-Centric Security

## LEVERAGING INDICATORS OF BEHAVIOR FOR EARLY DETECTION AND PREDICTIVE RESPONSE

# Introduction

In order for organizations to reach the largest portion of the market and deliver continuous value to their clients, product and service offerings must be able to scale instantly while remaining constantly accessible around-the-clock from almost anywhere in the world. Organizations are faced with solving for what has remained a Sisyphean task of maintaining a secure posture across each element of the IT infrastructure continuously.

To accomplish this, organizations need unparalleled visibility across a constantly evolving ecosystem and enhanced automation of security capabilities as part of the very DNA of the organization's security operations. Today's security model produces an endless stream of uncorrelated alerts for individual events on the network. The majority of these alerts are either false positives that need to be disqualified, or are simply glimpses of a larger attack sequence that will require an analyst to manually triage, investigate, and then correlate against other alerts—a process that simply cannot scale effectively to keep organizations and their clients secure.

There is a better way, but it requires a fundamental change in how we approach security by moving away from the labor intensive, inefficient and ineffective alert-centric model we continue to cling to in favor of a more effective, highly efficient Operation-Centric approach. Alerts are useful for identifying an element of an attack at a specific point in time, but they are far less effective in actually surfacing the entirety of an attack operation without a good deal of manual investigation and assessment.

For example, an alert may indicate the presence of a malware implant on a particular machine, where the machine can be isolated and the malware infection remediated—but successfully mitigating malware on one or more devices is equivalent to detecting and disrupting the entire attack operation, identifying root cause, hardening the infection vector, interrupting command and control (C2), eliminating persistence mechanisms, and so on. An alert-centric approach is akin to stopping the nosebleed while overlooking the brain tumor.

An Operation-Centric approach, on the other hand, seeks to quickly correlate all aspects of the attack sequence across disparate assets where event telemetry from each asset is largely subjective until it can be evaluated in the context of all available telemetry to arrive at an objective assessment that provides deeper visibility into an attacker's actions and activities.

An Operation-Centric approach can deliver detection and response automation at scale by leveraging Indicators of Behavior (IOBs), the more subtle signs of an attack that can surface the entire malicious operation at its earliest stages, allowing for earlier detections that inform a predictive response capability for comprehensive remediation that our current reliance on retrospective Indicators of Compromise (IOCs) can never deliver.

An Operation-Centric approach can deliver **detection and response automation at scale by leveraging IOBs,** the more subtle signs of an attack that can surface the entire malicious operation at its earliest stages.

# LOCUS OF CONTROL IS NOW DECENTRALIZED

In today's complex enterprise network ecosystems, there is no "hard perimeter" we can simply maintain to protect the "gooey center" where an organization's crown jewels reside. Many employees are now being onboarded for fully remote positions with little to no expectation they will revert to working in the traditional corporate office.

This means contractors, temporary workers, and regular employees are all now operating from home or from public locations while accessing the corporate network, pivoting to cloud-based productivity suites from on-premises network applications, and alternating between company-owned and their personal devices—all in the regular course of their daily work. Consequently, security operations have never been more challenged than they are today.

It is no longer just the IT department that is in control of introducing new hardware or software to the IT stack **—now anyone can make significant additions to the network environment in an ad-hoc manner** with little forethought or consideration of the potential security implications.

While IT environments had already been moving in the direction of this new work-from-anywhere paradigm, following the pandemic the rate of this shift has accelerated almost exponentially. With this shift we see a similarly rapid swing in the locus of control within the organization where governance of information systems is concerned. It is no longer just the IT department that is in control of introducing new hardware or software to the IT stack; now anyone can make significant additions to the network environment in an ad-hoc manner with little forethought or consideration of the potential security implications. This means that the security team is effectively no longer an extension of the IT apparatus where the two units work in tandem, and that many of the assets security teams must protect are not necessarily under the control of the IT team anymore.

Today's Defenders are no longer faced with just protecting computers, servers, file systems, and databases on-premises, they now need to protect an expansive, continuously evolving, heterogeneous fabric that is the modern distributed enterprise network. This is due in part to the fact that the majority of an organization's network isn't necessarily residing in local data centers anymore, but instead is dispersed across the networks of global cloud services and other SaaS providers, or are the domain of someone's potentially insecure home network.

In this new paradigm, compromise of these systems is inevitable, so the focus for security teams must be on automating more of their security operations to prevent the most common varieties of attacks, as well as implementing the right solutions to detect and respond to highly targeted malicious operations as early as possible. They also need to be able to reliably anticipate the emergence of new attack tools, techniques and procedures (TTPs) in a predictive manner so a mitigating response is ready when needed for novel attacks of the future. The adversary has already made a shift to capitalize on the "new normal" by expertly leveraging these diversification trends to their advantage, and Defenders are falling further and further behind in their ability to adequately address the threats.

The nexus for the security of all machines and devices, as well as the critical applications those devices are accessing and running either on-premises or in the cloud, requires a predictive detection and response capability that can correlate disparate elements within this complex and decentralized IT stack. In order to quickly identify any potential risk introduced into the environment, as well as any attempt to exploit potential vulnerabilities, the security apparatus must be fully automated and capable of proactively detecting potential compromise based on the most subtle chains of behaviors to reveal malicious activity before it can escalate to a major security event.

This is where the application of IOBs promises high-fidelity detections earlier as well as automation of response actions across the entire security stack. Here, the focus is placed on security value and actual outcomes instead of merely quantifying the frequency and category of events detected and blocked across the network based on the type of security tool deployed to monitor a specific asset type. Additionally, an Operation-Centric approach that is focused on detecting not just elements of an attack associated with individual alerts, but the whole of a malicious operation from root cause across every affected device, system, application, and user identity provides the basis for a predictive response capability where each action and activity the adversary takes works to reveal their intent and allows for automation of predictive response actions.

**The security apparatus must be fully automated and capable of proactively detecting potential compromise** based on the most subtle chains of behaviors to reveal malicious activity before it can escalate to a major security event.

# WHEN "MORE" DOES NOT EQUAL "BETTER"

Compared to other fields of specialization, the security industry is still a relative newcomer to the market. And while solutions like antivirus and firewalls have offered the illusion of a degree of automation in detecting and blocking threats, those capabilities are actually dependent on untold hours of ongoing manual signature development, manual triage, manual investigation, and manual mitigation efforts by humans. That worked relatively well for a short period when organizations were less frequently the target of attacks, but they do not translate well decades later when the adversary is automating attack sequences at scale across every potentially vulnerable asset.

Attackers quickly evolved to introduce tactics that rendered approaches like signature-based malware prevention largely ineffective against novel and highly targeted attacks. The adversary also began to automate aspects of their attacks to shorten the reconnaissance period and identify vulnerable systems more quickly, as well as automating other aspects like payload selection based on system specifications or in disqualifying targets based on language or region.

This is when we saw the number of attacks and new malware variants skyrocket. The simple fact is that manual human intervention simply cannot scale–the attackers know this and have automated where possible. The security industry followed suit by developing new classes of products that have increasingly employed more automated features so Defenders could try to keep pace. Unfortunately, when combined with an increase in solution specialization and a more complex security stack, Defenders soon found themselves overwhelmed by a flood of contextless, uncorrelated alerts which they still had to manually triage, investigate and respond to. That's when SIEM (Security Information and Event Management) solutions entered the scene.

SIEM solutions were intended to be one of the primary tools organizations could use to make sense of this avalanche of security and log data driving the flood of alerts. SIEM tools allow aggregation of telemetry from a variety of security tools and provide a centralized repository for security investigations and compliance mandates. In theory, SIEM solutions would employ a data lake structure and cloud analytics in an effort to normalize alerts and distill them down to the events that actually need an analyst's attention. But the value and effectiveness of a SIEM is highly dependent on the sources of data it has access to and how well it has been architected, tuned, and maintained. SIEM solutions require so much care and feeding that Defenders spend much more time managing and tuning their deployments than actually doing the job they were hired to do—mitigating threats.

When combined with an increase in solution specialization and a more complex security stack, **Defenders soon found themselves overwhelmed by a flood of contextless, uncorrelated alerts** which they still had to manually triage, investigate and respond to.

Unfortunately, SIEM solutions have an issue with returning too many false positives and they generate too many uncontextualized alerts, resulting in the very alert fatigue they were supposed to reduce, and as a result some high-priority threats end up being overlooked. SIEM tools can be useful for basic correlation, but they have not delivered on the promise of significantly reducing complexity in operation or in organizational risk, and they do not offer a means to coordinate an effective response for today's highly distributed architectures. In practice, analysts using SIEM solutions still have to manually sift through all the "well organized noise" in order to actually "find the signal," and then again they have to manually intervene to address the threat. That's where SOAR (Security Orchestration, Automation, and Response) solutions came on the scene, with a promise to leverage automation to streamline investigation and response.

As with SIEM, SOAR tools are only as good as their integrations with data sources, and users must be sure to test and retest frequently to minimize the risk of unintentionally impacting end-user experience or blocking critical processes. SOAR platforms also require skilled staff to not only manage the infrastructure, but to develop playbooks that often use Python or proprietary scripting languages—so again, security analysts spend more time managing and tuning the tool than they do actually using it. Ultimately, SOAR created a much more "organized mess," but it's all still a mess, so most organizations continue to look for new security approaches. While SOAR promised to usher in an age of automated response to threats, the level of manual human support required to deploy and maintain them has undermined their potential value.

**SIEM tools can be useful for basic correlation, but they have not delivered on the promise of significantly reducing complexity** in operation or in organizational risk, and they do not offer a means to coordinate an effective response for today's highly distributed architectures.

SOAR solutions were designed to extend beyond the use-cases of SIEM by providing a means to coordinate a response across a security stack comprised of multiple products. SOAR systems ingest and then analyze security event data to an extent similar to a SIEM, but go a step further by offering automated actions in response to specific events or triggers as defined by security policies. In theory, SOAR tools should automatically respond to low-level events or escalate alerts to notify security teams when human intervention is required. While SOAR tools help analysts by automating the gathering of data for investigations, they are generally not trusted to initiate automated responses without a human in the loop, particularly when dealing with critical business systems.

# THE SHIFT TO DETECTION AND RESPONSE

This brings us to the crux of the issue: our goal as the human element in the cybersecurity equation is to be able to readily ascertain the system state of our networks and extract the most complete story possible that will allow us to make sound decisions in order to mitigate organizational risk. As it stands today, we are still pretty far from realizing this goal because attackers are coming at us from the unobserved edges and seams within the larger IT ecosystem that are not covered by any one security solution, traversing between siloed assets and peripheral systems with relative ease, largely unconcerned about a timely response from Defenders.

Dwell time, the duration in which an attacker operates within a compromised network before they're detected, is probably one of the least talked about measures of a security program's efficacy. Various studies put the average dwell time at weeks to months to a year or more, which gives malicious actors more than enough time to conduct internal reconnaissance on a network, move laterally to access sensitive assets and escalate privileges, and then exfiltrate sensitive and proprietary information—activities that all drive up mitigation and recovery costs. This is where point solutions like Endpoint Detection and Response (EDR) really became game changers for the industry. Endpoints represent the largest segment of the addressable attack surface, so EDR provided the means to detect and disrupt intrusion attempts earlier, as well as provided the tools for Defenders to hunt for otherwise undetected threats in their environments.

EDR takes a step beyond traditional antivirus solutions by focusing on detection and response on an organization's endpoints, but it also has significant limitations. Often, malicious actors need to compromise a desktop, laptop, smartphone, server, or other endpoint to establish a foothold on a target's network, and they may need additional endpoints to move laterally and/or steal information. To defend against these malicious activities, EDR prioritizes continuous monitoring and threat detection as well as some level of automated threat response on each endpoint. Those two capabilities help analysts to quickly respond to endpoint threats. Indeed, EDR might be able to yield visibility into what's going on with an organization's endpoints, but the detection and response capabilities end there.

First of all, EDR can't detect threat activity on endpoints that don't have an EDR agent installed, so unless employees are willing to have agents installed on their personal devices, there is still a swath of unprotected endpoints connecting to an organization's network. Also, EDR lacks the visibility necessary to provide intelligence into how attackers might be combining infected endpoints with malicious activity in the cloud, or with compromised user identities, or across other parts of the network as part of a multi-stage operation.

**Various studies put the average dwell time at weeks to months to a year or more,** which gives malicious actors more than enough time to conduct internal reconnaissance on a network, move laterally to access sensitive assets and escalate privileges, and then exfiltrate sensitive and proprietary information.

However, attacks today are often more complex than just the detonation of a malicious executable or exploitation of a vulnerability on a single endpoint, so the scope of EDR is still too myopic to effectively defend against a broader malicious operation that includes actions and activities beyond the endpoint. This is where Extended Detection and Response (XDR) solutions come into play. XDR represents the first opportunity we have had to automate correlations across all asset types for earlier detection of malicious activity based on chains of behavior, as well as our first shot at developing a predictive response capability that takes the whole of an operation into consideration versus continuously responding to piecemeal alerts that only offer a fractional view of the attack.

XDR is in some respects a logical extension of the already widely adopted EDR approach, where EDR solutions offered a significant improvement over traditional signature-based solutions. But while XDR applies the same behavioral detection strategy that drives EDR, it significantly widens the scope of telemetry assessed to include all aspects of the modern distributed IT environment, including cloud workloads, applications suites, user identities, and more.

Attacks today are often more complex than just the detonation of a malicious executable or exploitation of a vulnerability on a single endpoint, so **the scope of EDR is still too myopic to effectively defend against a broader malicious operation** that includes actions and activities beyond the endpoint.

This more comprehensive approach to gathering and correlating telemetry from across the entire attack surface allows XDR to identify the more subtle patterns of malicious activity and detect potential threats earlier by "connecting the dots" across otherwise disparate assets. XDR works to break down the data silos that attackers rely on to remain undetected by unifying device, system, and identity telemetry for more efficient and effective threat detection and response by surfacing attacks earlier in the kill chain. The promise of XDR is even greater when the underlying conviction engines are driven by IOBs, which include seemingly benign activity that can be identified as malicious when the behaviors are chained—activity that would not trigger an alert by traditional security tools.

**The promise of XDR is even greater when the underlying conviction engines are driven by IOBs,** which include seemingly benign activity that can be identified as malicious when the behaviors are chained—activity that would not trigger an alert by traditional security tools.

# THE DIMINISHED VALUE OF IOCS FOR EARLY DETECTION

If we set aside untargeted commodity-based attacks and focus on the more complex and concerning operations, we know that IOCs are constantly changing and more often unique to a specific target, so leveraging IOCs for proactive defense in another environment is unlikely to result in earlier detections. Even the assumption that IOCs are somehow uniformly applicable in every instance for a given attack campaign in the same environment has proven to be demonstrably false. Furthermore, the more advanced attackers engaged with a high-value target often change their TTPs within the same kill chain when moving from one device to the next in a target environment, making early detection based on already known IOCs nearly impossible.

The more advanced attackers engaged with a high-value target **often change their TTPs within the same kill chain when moving from one device to the next in a target environment,** making early detection based on already known IOCs nearly impossible.

On top of that, attackers are keenly aware that a good deal of threat intelligence is shared within and between organizations, so when analysts are searching for specific IOCs in public repositories such as VirusTotal or similar, they are in effect telegraphing which techniques and tools have been flagged, and by omission which are still effective, because the attackers can also see the activity on the platform. This intelligence feedback is quite advantageous for attackers who then respond by tuning their TTPs accordingly.

To be clear, IOCs are still quite valuable for detecting known TTPs, just as outmoded signature-based detections are still effective for detecting common malware strains, and they will continue to be an important aspect of our security toolkits for the foreseeable future. But given the limitations for their application in surfacing highly targeted and novel attacks as described above, the question remains as to how we can detect more reliably and earlier in the kill chain. Most of the security efforts for the last two decades have been based on the idea that if we collect enough security telemetry in one place, we can just sift through it and find evidence of an attack—the SIEM/SOAR approach. But we still have not solved for the signal-to-noise ratio, and that's where IOBs can work to shift detections further to the left on the kill chain.

This means we need to focus not just on aggregating all available telemetry, but on also enriching the right telemetry to assure a high signal-to-noise ratio to detect advanced techniques used for initial ingress, persistence, privilege escalation, the compromise of user identities, and so on. And we need a means to consolidate all this enriched telemetry and operationalize it in defense of our networks in an automated fashion where response actions are initiated immediately. Furthermore, we need the ability to share this proactive intelligence with other security practitioners in the same way we are able to share retrospective IOCs today. Leveraging IOBs to achieve an Operation-Centric security posture represents the clearest path to achieving autonomous security operations at scale.

# DEFINING IOBs

IOBs describe the more subtle chains of malicious activity derived from enriched telemetry from across all network assets. Unlike backward-looking IOCs, IOBs offer a proactive means to leverage real-time telemetry to identify attack activity earlier, and they offer more longevity value than IOCs have ever been able to deliver. With the SIEM/SOAR approach, telemetry was aggregated and recorded, but if no standing policy was triggered, then all that telemetry is for the most part rendered ineffective. IOBs present the opportunity to maximize the value of the telemetry we already have.

IOBs describe the approach that malicious actors take over the course of an attack. They are based on chains of behavior that can reveal an attack at its earliest stages, which is why they are so powerful in detecting novel and highly targeted operations. Sooner or later, an attacker's path diverges from the paths of benign actors. But IOBs are not about just looking for anomalies or a key indicator of malice at a particular moment in time, although that's also part of it. IOBs are about highlighting the attacker's trajectory and intentions through analyzing chains of behaviors that, when examined together, are malicious and stand out from the background of benign behaviors on the network.

**IOBs are about highlighting the attacker's trajectory and intentions through analyzing chains of behaviors** that, when examined together, are malicious and stand out from the background of benign behaviors on the network.

IOBs can also be leveraged to detect the earliest signs of an attack in progress that are comprised of "normal activity" one would expect to see occurring on a network, such as we see with techniques like living off the land (LotL/LOLBin) attacks where legitimate tools, processes, and binaries native to the network are abused by the attacker. When otherwise "legitimate behaviors" are chained in certain sequences, they either produce an exceedingly rare condition or represent a distinct advantage for an attacker (or both)—this is where context-rich correlations across endpoints, the cloud, application suites, and user identities are critical for detecting malicious activity at the earliest stages of an attack.

When otherwise "legitimate behaviors" are chained in certain sequences, they either produce an exceedingly rare condition or represent a distinct advantage for an attacker (or both)—this is where **context-rich correlations across endpoints, the cloud, application suites, and user identities are critical for detecting malicious activity** at the earliest stages of an attack.

Instead of focusing on the episodes of the journey like IOCs, the utility of IOBs are independent of the attack sequence. Behavioral analysis has a chance of spotting processes that are behaving differently, sure, but such behavior only becomes clearly and demonstrably "bad" after several stages of the attack have been undertaken. Alternatively, IOBs can be leveraged to determine which "normal" behaviors are in fact malicious because the chains of associated behaviors themselves expose the attacker's intent even when the actions and activities fail to trigger an alert from traditional security solutions.

To better express the logic behind IOBs, let's take a look at the MITRE ATT&CK framework. The real value in the ATT&CK framework is that it creates a taxonomy for Defenders to better describe the actions and activities of the adversaries, leading to better detections. The columns in the MITRE ATT&CK represent the tactics employed in an attack, and each box represents the techniques used to accomplish them. Our job as Defenders is to ferret out the links between these tactics and techniques by ascertaining the trajectories, paths, and event sequences that stand out from the normal background noise of the network. This is achieved through first enriching the context of the individual telemetry sets, then by establishing correlations between the most meaningful telemetry to bring to the surface detectable moments early in the kill chain. The objective here is not to find one of the specific tactics listed in the ATT&CK framework, but instead to achieve an orthogonal view of an attack where we can determine the adversary's intended pathway by way of their actions and activities in a more predictive and proactive manner.

Every box in the MITRE ATT&CK framework represents a chance for Defenders to detect attack activity based on behaviors, and then to strengthen the correlation nodes that reveal the chain of associated behaviors that have come before, as well as predict the behaviors that are likely to come after. Behaviors that are ostensibly benign and within our expectations when observed in isolation become suspect when observed as chains of behavior because they are either statistically rare or present a distinct advantage to an attacker. Thus, detecting based on recognizing the utility of certain chains of behavior works to push detections further to the left and allows Defenders to respond in real-time to stop the malicious activity from progressing.

This requires the ability to observe and instrument both good and bad behaviors at scale and create data structures which enable queries that allow for additional context enrichment from a range of telemetry sources to answer the question 'are we under attack?' definitively" This does not imply advocacy for pursuing a shift to simple behavioral anomaly detections or going further down the UEBA path. It does mean that XDR solutions present the best available avenue for creating an extensible format for leveraging IOBs that can scale into the future as offensive tactics and our defensive capabilities continue to evolve.

The objective here is not to find one of the specific tactics listed in the ATT&CK framework, but instead to **achieve an orthogonal view of an attack where we can determine the adversary's intended pathway** by way of their actions and activities in a more predictive and proactive manner.

## SOLARWINDS ATTACKS

# A Case Study for Indicators of Behavior

Let's take a look at the role IOBs could have played knowing there were no IOCs available for detecting and blocking the novel attacks via SolarWinds where traditional defense approaches obviously failed to detect the threat. Malicious actors who really know what they're doing uniquely compile their code to make sure it doesn't match with any known file hashes or malware signatures out there, rendering IOCs largely ineffective for detection. These more advanced attackers also commonly inject false artifacts into IOC databases in order to ratchet up the noise and complicate response efforts. They also increasingly employ Living off the Land (LotL) techniques along with fileless malware in an attempt to leave as few traces of malicious activity behind as possible.

The malicious code used in the SolarWinds attacks was surreptitiously included in a legitimate software update signed by a valid digital certificate, making the attack extremely difficult to detect using traditional methods. In fact, it was only by sheer luck that an analyst at a third-party security vendor even noticed what was occuring. The takeaway here is in understanding that the key to early detection of advanced operations such as the SolarWinds attacks is in leveraging IOBs to level-up to a more efficient and effective Operation-Centric approach to detecting the whole of an attack as opposed to responding to individual, uncorrelated alerts that may or may not reveal key elements of the larger attack operation. Detecting and removing a malicious executable from an endpoint is a good thing, but it does not mean you have disrupted the larger offensive.

**In brief, some of the methods employed by the SolarWinds attackers included the following actions and activities:**

- The attackers had used the SolarWinds process to create an Image File Execution Options (IFEO) Debugger registry value for the process dllhost.ex. This technique enabled the attackers to establish persistence, but it also gave them the ability to execute malicious code when wscript.exe launched. That process ran a VBScript file, which in turn activated rundll32.exe. This executable invoked the Cobalt Strike DLL using a parent/child process tree that was separate from the SolarWinds process.
- The attack campaign prepared a unique Cobalt Strike DLL implant for each machine and avoided reusing the names for folders, files, and export functions, as well as other components like command and control (C2) domains and timestamps. In addition, it employed a domain generation algorithm (DGA) to evade security countermeasures.
- The attackers renamed their tools and binaries and placed them into folders that impersonated legitimate programs native to the network and files that already existed on the compromised machines.
- Prior to engaging in extensive keyboard activity, the threat actors used AUDITPOL to disable event logging and used the same tool to re-enable it after they were done.
- The attackers similarly prepared special firewall rules to limit outgoing packets for certain protocols while they performed their espionage activities on a targeted network. They then removed those firewall rules once they had concluded the reconnaissance stage of their operation.
- The attackers carefully planned for each instance of lateral movement by first enumerating remote processes and services that were running on the targeted host. They then moved laterally only after they had disabled some security services.
- The attackers used timestamping techniques to change the timestamps of artifacts along with a sophisticated wiping process to make it more difficult for investigators to find and recover their DLL implants.

All of this activity and much more went undetected by alert-centric security approaches that are dependent on IOCs and traditional threat intelligence sources but lack the ability to enrich the context around individual events and chain them through automated correlation with other behaviors occurring on the network. Yet, there was no shortage of detectable activity in the environment if the security solutions in place were capable of recognizing the IOBs:

- **Attackers used the SolarWinds update process to spawn a separate process that invoked the Cobalt Strike DLL:**
  A trusted process used for digitally signed software updates typically would not spawn a separate process unassociated with the product in question, so this falls in the "rare behavior" category that should elicit suspicions, especially if the spawned process invokes a DLL stager not necessary for the trusted software update. While these behaviors are not outwardly malicious in and of themselves, aside from being rare they also fall into the "advantageous behavior" category when chained together. Detection of a malicious operation should not be undermined by the system "trusting" what is assumed to be "legitimate" files or processes. If the behavior of those files or processes is rare or of potentially high value to an attacker, they can be used to reveal the attacker's intentions, even when there is no outward sign that the inherent "trust" in the principal file or process has been violated.

- **Attackers renamed tools and binaries, employed DGA for obfuscation:**
  The use of unique Cobalt Strike DLL implants for each target machine, avoiding nomenclature reuse for folders/files, and the use of a DGA to mask C2 communications all render IOCs and other rule/signature-based approaches ineffective for detection. But IOBs can be leveraged here for binary similarity analysis where instead of looking for the "fingerprint" of "known" malware, the DNA of an executable's coded behavior is matched to comparable behaviors of malware for detection of even novel variants. As well, IOBs can be leveraged in packet capture cluster analysis to identify methodologies designed to bypass traditional security approaches like use of DGA to mask C2 communications. The tools and tactics may be obfuscated by the attackers, but when the behaviors are chained they reveal attacker trajectories that stand out from the background noise of "normal" network activity.

- **Use of AUDITPOL to disable event logging and the addition/removal of firewall rules:**
  AUDITPOL is a "legitimate" tool used to configure and manage audit policy settings and logging. Similarly, reconfiguration of firewall rules is something that normally occurs on a network from time to time. When these behaviors are chained together, or chained with other activities that require elevated user privileges, they stand out in the context of rarely observed events that would produce optimal outcomes for an attacker to reveal key actions early in an attack sequence, even in the absence of an alert.

- **Lateral movement on the network:**
  Lateral movement further colors the whole behavioral chain up to this point by providing context for the earlier behaviors where the unit of escalation and subsequent interactions is the behavior chain that reveals the attacker's intent, where an alert merely points to a moment in time for the longer chain of events in the attack sequence. It is here that the utility of IOBs really shifts from retroactive detection of activity that has happened or is happening now, to one of predictive value in forecasting the appropriate response action to prevent the attacker behaviors that would logically follow. Here, chains of otherwise disparate behaviors have coalesced to reveal the attacker's pathway and likely next steps, allowing for predictive, proactive, and automated response actions that anticipate the trajectory of the kill chain and disrupt the entire malicious operation, not merely an aspect or element of the attack.

- **Timestamping of artifacts:**
  The utility of IOBs are not dependent on retrospective forensic investigation or any evidence examined after an attack has occured. While timestamping will certainly work to hamper forensic investigation for compliance and reporting purposes, the technique does nothing to hamper detection of the attack based on the chains of behaviors described above.

# THE APPLIED UTILITY OF IOBs

Leveraging IOBs can empower security operations to more quickly identify and understand the full scope of threats by correlating suspect actions and activities. This allows analysts to focus on managing security policies that drive automated mitigation actions to assure they are consistently applied across the entire network regardless of its size or complexity, as opposed to analysts being endlessly consumed with triaging and investigating individual, uncorrelated alerts.

So how does detecting based on certain chains of behavior accomplish this? Let's take a look at the logical stages of chained behavioral detections and how they can catalyze effective automation of the entire defense process from monitoring and telemetry enrichment to detection and remediation.

Leveraging IOBs can empower security operations to **more quickly identify and understand the full scope of threats** by correlating suspect actions and activities.

## STAGE 1 ▸ ESTABLISHING A COMMON EXTENSIBLE LANGUAGE

Early detections based on chains of behavior will require different security solutions to have the ability to deliver telemetry in a common, extensible language format with a shared nomenclature. For example, take the case of end users and their email accounts where a single user may have multiple user identities, IP addresses, and devices that are all generating their own telemetry that is processed in completely different ways and is being analyzed in distinct siloes. This situation makes it much more difficult to determine if any or all of these assets are being abused in the earliest stages of an attack without a good deal of manual investigation.

This siloed structure, and the likelihood there are differing telemetry formats that are not comparable across all network assets, presents a significant obstacle to automating correlations across all of the available telemetry to surface a complex attack early, or even at all. The lengthy manual investigation cycles required to overcome these obstacles presents a resource constraint issue for organizations and precludes any advantages that automation can offer because human investigations simply do not scale.

Contrast that scenario with the case where a human Defender is tasked with analyzing telemetry from a unified data structure where all the intelligence generated by an array of solutions has been normalized by a common format and means of expression. In this situation, an analyst can more easily understand the correlations between an event on an endpoint, for example, and how it relates to potentially compromised credentials used to login to a SaaS solution and any subsequent activity like accessing sensitive data or attempting to escalate privileges.

The analyst can see the correlations across otherwise disparate assets and events to more readily make an assertion as to whether these chains of behavior are malicious, as well as initiate remediation actions per the organization's security policies. While this is a better situation with regard to detection and remediation, it is still too dependent on manual human processes which, again, simply cannot scale—but once we have the telemetry organized in a common, extensible language and format we are much closer to a process that can.

## STAGE 2 ▸ CODIFYING CHAINS OF BEHAVIOR ACROSS MULTIPLE ASSETS

Most systems are collectors of security and other related telemetry, with only the occasional behavior revealed in the process. What's needed is more behavioral instrumentation of the available telemetry to gain the capability to recognize chains of behavior that may be malicious in order to detect stealthy attacks such as with an operation like the SolarWinds campaign. Even when attackers are doing things that could be regarded as normal activity that one would expect to see on a network, when the actions and activities are chained, a pattern emerges that is divergent from the behavioral paths of authorized users and activities.

By leveraging IOBs as a means to codify both normal and abnormal chains of behavior, it is possible to not only quickly gain full visibility into an attack that's already happened (as with IOCs), but it is also possible to use that same progression of behaviors to recognize novel attacks now and into in the future. IOBs are not about anomaly hunting or searching for a key indicator of malice, although that is a lesser part of their utility. Leveraging IOBs is about highlighting attacker pathways via chains of behaviors that stand out from the background noise of all the other behaviors occurring on the network. This allows us to build a robust detection layer on top of a unified data structure that can not only ferret out likely attack scenarios in progress, but also form the foundation for being able to predict potential attack pathways so they are quickly recognized and terminated across the network, or on other networks in much the same way that IOCs are currently employed today.

When all associated behaviors are considered as a singular operation, what emerges are just two possible sets of behaviors: **those that are clearly benign and those that are clearly malicious.**

## STAGE 3 ▸ AUTOMATING BEHAVIORAL DETECTIONS AT SCALE

Once we have the ability to map telemetry data of both good and bad actions and activities behaviorally, we can leverage the application AI/ML to make the necessary distinction between the two categories, to correlate otherwise weak signals of an attack earlier in the kill chain, and then ultimately assume a predictive response stance in addressing risk mitigation. Key here is understanding that the types of activities that can occur on a network are ultimately finite, and the possible combinations of those activities that have utility are even fewer. Thus there are a finite number of definable actions that an attacker can engage in to successfully execute an attack. When all associated behaviors are considered as a singular operation, what emerges are just two possible sets of behaviors: those that are clearly benign and those that are clearly malicious.

This is critical in defending against attackers who are constantly developing new techniques that evade IOC-based defenses. This is about having an Operation-Centric strategy in place that understands attacker TTPs as a subset of all possible user, system, and device behaviors to gain the visibility into complex attacks that will enable Defenders to detect and block threats at the earliest stages and then eventually lead to a state of autonomous security operations. As with the SolarWinds supply chain attacks, threat actors will continue to evolve their approach to assure they evade defenses and blend in, but they will not be able to mask the malevolent nature of the behaviors when analyzed in combination operationally.

AI/ML automation provides an obvious advantage here in providing the means to ascertain the behavior chains of interest while disregarding those that need no attention, and to do it consistently at scale. As well, the application of AI/ML also delivers the capability to reveal previously unrecognized chains of potentially malicious behavior as TTPs evolve and change, in addition to predicting new attack pathways as our IT and security stacks evolve and change with the addition of new technologies. The outcome here is the delivery of a single dynamic graphical representation of these chains of behavior that reveals all of the possible relationships between users, devices, assets, and other associated elements abused in an attack.

This is where the "organized mess" that is the telemetry aggregated by tools like SIEM and other security solutions that warehouse uncorrelated events and alerts can add valuable color to the detected chains of behavior. They can inform the detections by providing the specific user, device, and other intelligence that illustrates the entire attack progression in detail from root cause across all affected assets.

AI/ML can also be used to automate the appropriate responses to mitigate events faster and more completely, as well as make a determination when and if a chain of behaviors needs to be escalated to a human in order to evaluate the efficacy of the security policies governing the entire detect and respond cycle. These behavioral detections remain evergreen in that they are more or less immune to new obfuscation techniques or to changing TTPs because—as mentioned above—the number of actions and activities an attacker can perform on a network are finite, so the utility of IOBs means it's actually the attackers themselves that work to surface their own malicious operations.

The number of actions and activities an attacker can perform on a network are finite, so the **utility of IOBs means it's actually the attackers themselves that work to surface their own malicious operations.**

## STAGE 4 ▸ FULLY AUTOMATED PREVENTION, DETECTION AND RESPONSE

At this stage, we now have the prerequisite capabilities to build out completely autonomous workflows that will essentially eliminate the risk from both known and unknown attacks by significantly raising the threshold for success for attackers. This is the stage where we move from automation of some security tasks to a state of fully automated prevention, detection and response with machines handling nearly every aspect of security operations, while humans focus on governance and administering policies based on an organization's risk tolerance and requirements.

## STAGE 5 ▸ THE DEMOCRATIZATION OF SECURITY

This Operation-Centric approach presents an opportunity for Defenders to share high-fidelity detections with other Defenders that are effective in any environment with any combination of hardware and software. This where the larger community of Defenders can collectively create a unified phalanx to confront the attackers where they operate, in real-time and at scale, and usher in an age of democratized security where no organizations are at risk due to a lack of resources or capabilities.

This ability to autonomously and proactively end attacks moves organizations to a sustainable "prevent" posture that would revolutionize security from the bottom up by relieving Defenders of the repetitive triage, analyze, and respond cycles so they can instead act as an "iterative tuner" of the security machine.

# LEVERAGING IOBs TO ACHIEVE OPERATION-CENTRIC SECURITY

An Operation-Centric approach reorients the detection and response cycle to focus on comprehensively disrupting entire attack progressions from root cause to every affected asset, versus the current focus on alerting on and remediating individual elements of the larger attack campaign. In leveraging IOBs, an Operation-Centric approach also presents the opportunity to create a repository of detectable behavior chains that can surface even the most novel of attacks as well as automated response playbooks to disrupt attacks at their onset. Understanding attacker intentions and likely pathways based on early-stage actions and activities enables defenders to proactively predict and disrupt subsequent stages of an attack, as well as provides an avenue to develop fully autonomous security operations.

In order to achieve a truly Operation-Centric posture and move closer to autonomous security operations, a future-ready standard that universally defines and operationalizes IOBs to ascertain the the actions and activities of the attackers is required. To be truly useful, there needs to be a common definition, language, and expression of IOBs that is completely independent of any particular security tool or vendor. The wide array of solutions available can provide the raw telemetry, as well as the color and context required to collectively interpret observable behaviors. But, as it stands today, security tools themselves don't provide a standardized language that can accurately describe and operationalize the chains of behavior that will enable us to detect and respond to attacks faster than the adversary can adapt.

Understanding attacker intentions and likely pathways based on early-stage actions and activities **enables defenders to proactively predict and disrupt subsequent stages of an attack,** as well as an avenue to develop fully autonomous security operations.

An Operation-Centric approach also presents the opportunity to leverage AI and ML to automate not only detections but also the majority of remediation actions that should follow. This would ultimately elevate the human analyst to focusing on tuning the policies driving detection and response, where they would be working on understanding possible attack sequences based on attacker intent in the context of critical business operations and the steps that should be taken to reduce overall risk to improve security operation efficiency and efficacy.

We need a common, extensible format for IOBs that can keep us all on the same page, yet is capable of scaling as our capabilities and those of our adversaries continue to evolve. This does not mean simple anomaly detection, shifting our focus to some sort of behavioral anomaly detection, or launching some sort of UEBA redux. Instead, we're talking about instrumenting and collecting both good and bad behavior at scale regardless of solution or asset class—and codifying the data structures that will enable automation of machine-language queries.

Operationalizing IOBs will require standardization that will deliver the full potential value of the entire security stack to quickly and autonomously deliver the necessary context and correlations across diverse telemetry sources without requiring that the raw telemetry first be "filtered" in order for the volumes of available data to be more manageable. Only recently have we seen the advent of big data platforms that have the ability to ingest and process the terabytes of information required to make autonomous detection and response at scale a possibility.

This is critical for defending where attackers develop new techniques that evade IOC-based defenses. This is about having a language in place to understand TTPs as a subset of all user, machine, and object-coded behaviors. Now, more than ever, we have the opportunity to focus on innovation and advancing the development of more effective detection capabilities rather than just deploying different iterations of the same old toolkit. We need to look beyond tools that only leverage retrospective IOCs, since this artifact-based approach obviously failed to detect the SolarWinds campaign and similarly complex attack sequences.

Instead, we need to look to IOBs that describe the rare and advantageous chains of behavior that created the foundation for campaigns like the SolarWinds attacks to be so successful. More specifically, we need to shift away from our reliance on artifacts that only describe what already happened and move towards leveraging the Indicators of Behavior that reveal what is happening in real-time in order to predictively respond to prevent the next steps of an attack progression long before it can escalate to a major security event.

An Operation-Centric approach that leverages IOBs empowers security operations to dynamically adapt and predictively respond more swiftly than attackers can modify and adjust their tactics to circumvent defenses, which is key to finally reversing the adversary advantage and returning the high ground to the Defenders.

We need to shift away from our reliance on artifacts that only describe what already happened and **move towards leveraging the IOBs that reveal what is happening in real-time** in order to predictively respond to prevent the next steps of an attack progression long before it can escalate to a major security event.

# ABOUT CYBEREASON

The Cybereason Data Science, Threat Research and Engineering teams are pioneering the study and application of Indicators of Behavior (IOBs) for faster, more comprehensive and more readily actionable detections that inform and drive predictive response capabilities and reduce risk for organizations by identifying and disrupting attacks earlier in the the kill chain.

The Cybereason MalOp™ leverages Indicators of Behavior, the subtle chains of behavior that reveal attacker actions and activities at the earliest stages of an attack, to generate suspicions that are enriched through contextual correlations that generate the MalOp (malicious operation) detection.

The MalOp delivers an interactive, in-memory graphical view of the full narrative of an attack sequence from root cause correlated across all impacted devices, systems, applications and user personas in a single screen. This unparalleled automated analysis increases detection speed and accuracy by reducing the noise of alerts with a focused deconstruction of the overall operation with all the information an analyst needs to scope and respond to a malicious operation to reduce the Mean Time to Respond (MTTR).

**Every MalOp generated includes the following five critical categories of information about a malicious operation:**

- **Comprehensive Detections from Root Cause:**
  All of the malicious activity that caused Cybereason to suspect that a malicious operation might be taking place, whether through a well-crafted spear phishing email or other entry vector, cyber adversaries must first establish a foothold in the environment to build on and escalate the intrusion. The root cause along with any other suspicious behavior is mapped to the MITRE ATT&CK framework.
- **Correlating All Impacted Users and Machines:**
  Although a specific user or asset might be the ultimate target, multiple systems will be leveraged along the path to their objective. All of the users and machines that are part of this larger operation are correlated into this single attack sequence view so analysts can easily determine the full scope of the operation and the breadth of compromise to drive a thorough and comprehensive response.
- **Identifying all Incoming and Outgoing Communications:**
  Data exfiltration and command and control (C2) activity are excellent beacons to further understand an attacker's intent, which informs a predictive response that disrupts the entire operation. Incoming and outgoing network traffic across all impacted machines is provided and traffic identified as malicious is highlighted for further analysis.
- **Tactics, Techniques, and Procedures (TTPs):** Quickly understand what TTPs the attacker is using to execute their malicious code and traverse the environment, including stealth and persistence mechanisms, as well as any Living Off The Land (LOL) techniques abusing legitimate assets that allow the attackers to hide their activity amongst otherwise normal network activity. All TTPs are mapped to the MITRE ATT&CK framework.
- **Full Timeline of the Attack:**
  The Cybereason MalOps automatically analyzes all attacker activity across the entire environment to automatically deliver the full timeline of the attack progression visually, giving analysts the ability to quickly drill down into each element of the attack and implement automated or one-click remediation so they can move beyond triaging and investigating individual alerts and leverage fully contextualized and correlated attack stories in real-time without the need for complex queries and protracted investigations.
- **Leverage All Available Telemetry:**
  Other solutions are forced to severely limit the amount of critical data collected and analyzed because their platforms simply can't process the huge volumes of data that are available, but Cybereason ingests 100% of event data in real-time and analyzes more than 9.8PB of threat intelligence weekly. Cybereason reduces investigation periods by as much as 93% so defenders can eliminate threats in a matter of minutes rather than days, elevating Level 1 and 2 analyst capabilities to Level 3 proficiency and delivering an unmatched 1:200,000 analyst-to-endpoint ratio.

Cybereason is *the* XDR company, partnering with Defenders to end attacks at the endpoint, in the cloud, and across the entire enterprise ecosystem. Only the AI-driven Cybereason Defense Platform provides planetary-scale data ingestion, operation-centric MalOp™ detections, and predictive response that is undefeated against modern ransomware and advanced attack techniques. Cybereason is a privately held international company headquartered in Boston with customers in more than 40 countries.