

®



4 Keys to Automating Threat Detection, Threat Hunting and Response

Contents

3	Executive Summary
4	Maturing Advanced Threat Defense
5	4 Must-Do's for Advanced Threat Defense
	1. Use New Post-Breach Detection Techniques to Reduce Dwell Time
	2. Automate Detection and Response Processes
	3. Evolve Capabilities to Proactively Hunt for Threats
	4. Engage Managed Detection and Response Services
10	Fidelis Elevate™ – Your Force Multiplier for Automating Detection and Response
12	Summary

Executive Summary

Cyber attacks are no longer unexpected. In fact, just the opposite. Incidents proliferate as threat actors enjoy continued success with phishing and web drive-by attacks. Ever evolving, attackers often shift their approach to file-less techniques using macros and scripts to evade preventive defenses and remain undetectable. Social engineering and business compromise scenarios, no longer the amateurish efforts of yesterday, operate outside the scope of defenses and play on human urgency, expectations, and trusted entities. Not to be forgotten are nation-states and organized crime. Their long-term reconnaissance, quiet entry, and ability to remain below the radar and maintain persistence within targets continue to set a high bar for sophisticated cyber attacks.

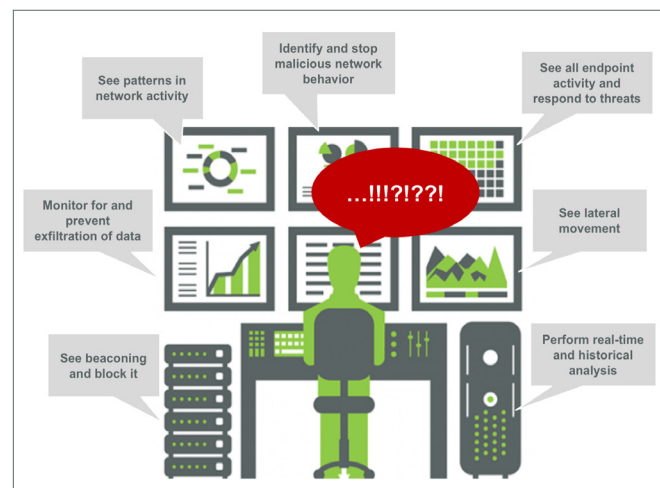
While the mindset of security leaders has shifted from keeping bad actors and malware out to realizing that malicious intruders and insiders are operating within environments undetected, organizations remain ill-prepared and hampered in their efforts to enact post-breach detection and response efforts.

As attackers continue to succeed, security leaders have responded by spending millions of dollars to consolidate alerts, events, and logs into SIEMs with little to no improvement in post-breach attack detection or reduction of dwell time. Despite investments in preventive technologies, attackers routinely compromise seemingly secure organizations and steal financial assets, intellectual property, and sensitive data.

Rather than help, preventive technologies hinder post-breach detection efforts as they are often noisy and generate multitudes of innocuous alerts that lead to alert fatigue. Alerts multiply as the same artifact or behaviors are detected at different stages of the attack lifecycle and duplication of alerts further adds to alert fatigue. More problematic, such technologies deliver neither the visibility nor the rich metadata necessary to detect and respond to attackers already within networks. Alerts generated by legacy security infrastructure lack sufficient contextual information and cross-session analysis to enable a security analyst to piece together the story from multiple point products, each alerting on different aspects of the attack. Because the current situation lacks a common metadata model, automation is difficult to apply. Without automation that would gather data and speed triage and investigation, security analysts must validate events while gathering and analyzing information from multiple disparate systems that neither integrate

nor share data well. Rather than being on their toes, out in front of the attack, security analysts must piece together evidence and clues and investigate manually – often on their heels.

Beyond the alert overload, lack of visibility, and insufficient contextual analysis, one of the key issues in cybersecurity is the skills shortage. Not only are too many disparate tools unused because of the lack of skilled staff to operate them, but the shortage of cybersecurity analysts skilled in post-breach detection and investigation is leaving gaps in coverage, stalling investigations. Even if enterprises could find the right talent, many face restricted budgets to expand security operations teams.



The lack of proper tools combined with alert overload, alert fatigue, and manual processes creates a perfect storm wherein the most important alerts supported by content and context within larger conclusions may be missed or, because of limited staff, dismissed. Even more

worrisome, the chances are high that hidden threats are already in the organization's network. Focused human adversaries know how to evade most security and monitoring tools by making their attacks look like normal activity. While organizations want to hunt for unknown threats, the state of being 'constantly behind' makes it virtually impossible to get in front of the situation – even if the tools, data, and expertise were available. It's well

understood that the more accurate the post-breach detection, the faster security teams will be able to address the risk; however, with limited visibility, security teams are unable to detect blind spots and black holes and are forced to act on incomplete or, worse, inaccurate information – leaving organizations open to increased risk and potential data loss or theft.

Maturing Advanced Threat Defense

When it comes to detecting and responding to threats and preventing data loss, speed and accuracy are everything as the machine initially compromised is almost never the one the intruder needs to accomplish his or her objective. The adversary must move laterally, burrowing deep in the network to find desired targets. With an average breakout time of less than two hours from a compromised system to lateral movement, security teams need to be able to quickly detect and respond to attacks that penetrate security defenses and resolve them before they do irreparable harm. However, the gap between time of compromise and time of detection is one of the main failures noted when investigating a breach.

With threat actors claiming victim after victim, it's obvious that an outdated, fragmented approach to cyber defense is unsustainable. To defend against determined attacks, organizations must mature advanced threat defense capabilities in order to reduce the Mean Time to Detect (MTTD) and apply automation to improve their Mean Time to Respond (MTTR). To that end, industry analyst firm Gartner recommends that IT security leaders shift from prevention-focused defense to one that prioritizes post-breach detection and response and lays the groundwork for threat hunting. In order to do so, Gartner recommends adopting a curated technology

stack from vendors that integrate Network Traffic Analysis (NTA), Endpoint Detection and Response (EDR), and sandboxing; and to consider lean-forward defenses such as deception. In addition to evaluating vendors on their ability to deliver a curated security technology platform, Gartner recommends that enterprises also evaluate vendors on their ability to provide Managed Detection and Response (MDR) services using the same curated technologies.

By adopting an integrated technology stack, organizations can:

- Increase visibility and detection accuracy across networks, clouds, endpoints, and enterprise IoT devices
- Confirm and stop data theft by inspecting the content of all outgoing network traffic
- Achieve deeper insights through consolidated and correlated information that delivers content and context for post-breach detection and response
- Strengthen the security stance by identifying and responding to attacks throughout the attack lifecycle
- Maintain in-house knowledge and control over security technologies for optimal use and advancement of security capabilities

Did You Know?

The gap between time of compromise and time of detection is one of the main failures noted when investigating a breach.

4 Must-Do's for Advanced Threat Defense

Beyond evolving the security infrastructure to a single network and endpoint detection and response platform, Gartner recommends that organizations evolve the post-breach detection and response capabilities of the security operations team. Four must do's for maturing advanced threat defense include:

1. **Use new post-breach detection techniques to reduce dwell time or MTTD**
2. **Automate detection and response processes to improve MTTR**
3. **Evolve capabilities to proactively hunt for threats**
4. **Engage managed detection and response services from technology vendors with in-house knowledge of their solution, plus the ability to provide incident response services**

1 Use New Post-Breach Detection Techniques to Reduce Dwell Time

The longer an attacker stays on your network the more damage they could do; therefore, it's imperative to reduce dwell time by using new post-breach detection methods to detect threats in the network.

Combined styles of threat defense. Taking guidance from Gartner's Five Styles of Advanced Threat Defense Framework, security teams are advised to use an updated defense model that focuses on two dimensions of post-breach detection. The first dimension, *where to look*, leverages network traffic analysis, endpoint detection and response, and payload analysis using a sandbox environment to identify threat activity and compromised endpoints. The second dimension, *when to look*, considers the real-time/near real-time or post-breach timeframe of an attack.

According to Gartner, the most effective post-breach detection comes from looking for activity in different parts of your network, both in real-time or near real-time as well as performing post-compromise forensic analysis. The five styles of advanced threat defense are:

- Network traffic analysis (network: real-time/near real-time)
- Payload analysis [sandboxing] (payload: real-time/near real-time)

- Endpoint behavior analysis (endpoint: real-time/near real-time)
- Network forensics (network: post compromise)
- Endpoint forensics (endpoint: post compromise)

The framework also highlights which combinations of styles are the most complementary. Effective post-breach detection comes from combining different styles (for example: network/payload, payload/endpoint, or network/endpoint). Because of the challenges in combating targeted attacks and malware, security-conscious organizations should plan on implementing at least two styles, the most effective of which are network traffic analysis (NTA) combined with endpoint forensic (EDR) and payload analysis combined with endpoint forensics (EDR).

Deception for post-breach detection. In addition to advising organizations to leverage the Five Styles of Advanced Threat Defense framework to lean forward in their post-breach detection and response capabilities, Gartner also recommends the use of deception technologies to better detect attackers that have penetrated their defenses.

We are all familiar with the old paradigm that an attacker just needs to get in once, and that security needs to do everything right, every day, every time to defend. The use of deception technology introduces a new paradigm,


wherein the attacker must do everything right to avoid detection and prevent digital 'noise' that would reveal their location and activity.

Cyber deception, recently identified by Gartner as a top security technology, can more effectively detect attacks that have infiltrated an organization's network, divert the attacker, and provide understanding around what assets have been compromised. By discovering and classifying network assets, you can easily create a decoy environment that is always current with your real environment and which adapts as your network assets change. As a result, the use of decoys, traps, lures, and other mechanisms is quickly gaining the attention of organizations seeking an efficient and highly accurate method for detecting human and malware attacks.

Deception technology shines a detection light on enterprise IoT and other assets where agents cannot be deployed. It slows attackers by deceiving them into believing they're interacting with real endpoints, servers, and data repositories; and prevents further spread of the attack by deflecting it from real assets to decoy assets. Deception improves and becomes deterministic with breadcrumbs on real assets, luring attackers, malicious insiders, and automated malware to decoys.

By taking an automated approach to deception deployment and maintenance, organizations alter the playing field for attackers. Instead of searching in vain for the bad actor within an

ocean of good data, deception delivers validated alerts and events from decoys, MITM (Man in the Middle) traps, Active Directory breadcrumbs, and traffic analysis. The threat defense framework and deception equip security operations teams to work with exceptional effectiveness and efficiency and to move from alerts to conclusions.



2 Automate Detection and Response Processes

Modern attacks are a complex, often-automated series of processes, steps, and interrelated events that penetrate the cybersecurity perimeter and generate significant cost, noise, and anxiety for an enterprise. Enterprises incur these challenges because the typical short-staffed security operations team lacks complete, integrated, and automated technology to detect, prevent, or respond to these attacks.

Automated detection and response can help organizations address the skills shortage by increasing the capabilities of a security team. It acts as a force multiplier to increase a team's capacity as well as the level to which it's able to operate. By automating data collection, investigation, prevention, and response processes within the security operations workflow, organizations can maximize a team's efficiency throughout the entire lifecycle of an attack.



Automate data collection. When it comes to investigating threat activity, it all comes down to data – do you have enough and is it the right type of data. Because incomplete, incorrect, or insufficient data can hamper investigations, it's vital that security analysts have the right data readily available. That means automating data collection.

Capturing data from all sensors all the time is futile and costly. Such packet capture (PCAP) data is high volume, expensive to store, and almost impossible to search and query or apply threat intelligence to. Metadata, on the other hand, is lightweight, much less expensive to store, and – because it is indexed – easy to query and search. One can also use threat intelligence with metadata for threat detection or threat hunting. Further, solid metadata with attributes of high variability are required for machine learning anomaly detection and various use cases. Possible network-sensor locations to capture metadata include direct gateways, internal networks, cloud VMs and VTAPS, email/MTAs, and web gateways. Endpoints can also capture metadata, and optionally capture memory, files or full disk images for forensic evidence, either through automation or on-demand.



Automate investigation.

Automating investigations can significantly reduce the alert fatigue caused by noisy legacy threat detection systems. Automated alert prioritization ensures the team focuses on alerts associated with the highest level of risk. Intelligent grouping of related alerts under the relevant, highest priority alert lets security teams focus on those that really matter and respond faster to interrupt the attack process. By providing a deeper level of insight, such as a sandbox execution report or all endpoint activity before and after the event, actionable alerts can elevate tier 1 or junior analysts to the level of tier 2 analysts, empowering them to quickly draw conclusions and rapidly respond to validated alerts to complete investigations.



Automate prevention.

A third area to automate is prevention. Automated prevention workflows can reduce the impact on a busy, overwhelmed security operations team, ultimately speeding up the prevention of new and previously unknown threats. Effective ways to reduce workload through automation include dropping sessions, redirecting web pages, quarantining suspicious/malicious files, blocking file execution, and preventing process execution on endpoints.



Automate response.

With automated response, an organization can considerably accelerate its response time. By setting up rules and custom scripts to automate endpoint response, such as isolating a compromised endpoint, wiping a file, or restoring from a backup image if a specific alert rule is triggered, analysts can both reduce the risk associated with specific alerts as well as save time by reducing the need to manually initiate a workflow for each individual alert.



Automate deception.

Deception defenses with automated discovery can learn a new environment to profile assets, resources, services, applications, and activity – as well as unknown legacy systems and shadow IT. From this rich profile base, auto-generated decoys can be created to show where and how deception layers would be used. However, creating and deploying deception layers is only one-

half of the effort – maintaining, adapting, and changing decoys and breadcrumbs for freshness is the second half. For deception to be effective, it needs to mirror the current environment, show daily activity with logins for fake AD accounts to decoys, and change regularly to maintain the freshness to avoid fingerprinting by attackers. Deception defenses can automatically adapt to changes in networking and resources given they keep a rich profile from automated discovery used during decoy creation. Automated adaptation of deception defenses is preferable to keep deception layers as realistic as possible and minimizes manual effort.

3 Evolve Capabilities to Proactively Hunt for Threats

The fact that dwell time between an attack and detection can take months begs the question: “Why not hunt for threats?” A just question, but the answer is not so simple. For many organizations, threat hunting is still a new and poorly understood process. Most organizations, believing themselves to be hunting threats, are still performing threat detection and responding to alerts. In order to mature advanced threat defense, it’s therefore important to understand the difference between threat detection, threat modeling, and threat hunting.

- Threat Detection leverages multiple detection techniques from signatures, rules, and patterns to anomaly detection, machine learning, and behavioral analysis to find known threats, query, or model. For example, matching indicators of compromise to various data sources is a form of threat detection, or searching through a SIEM or a collection of logs.
- Threat Modeling is a proactive process to improve applications, systems, and network security by assessing potential risks, threats, and vulnerabilities often from an attacker’s perspective, and then prioritizing countermeasures to address the effects.
- Threat Hunting is a proactive, analyst-centric, iterative, and interactive ad hoc process driven by expert intuitive hypotheses assuming a breach. The practice combines security expertise, data analyst skills and creative thinking upon a knowledgebase across applications, systems, and networks.

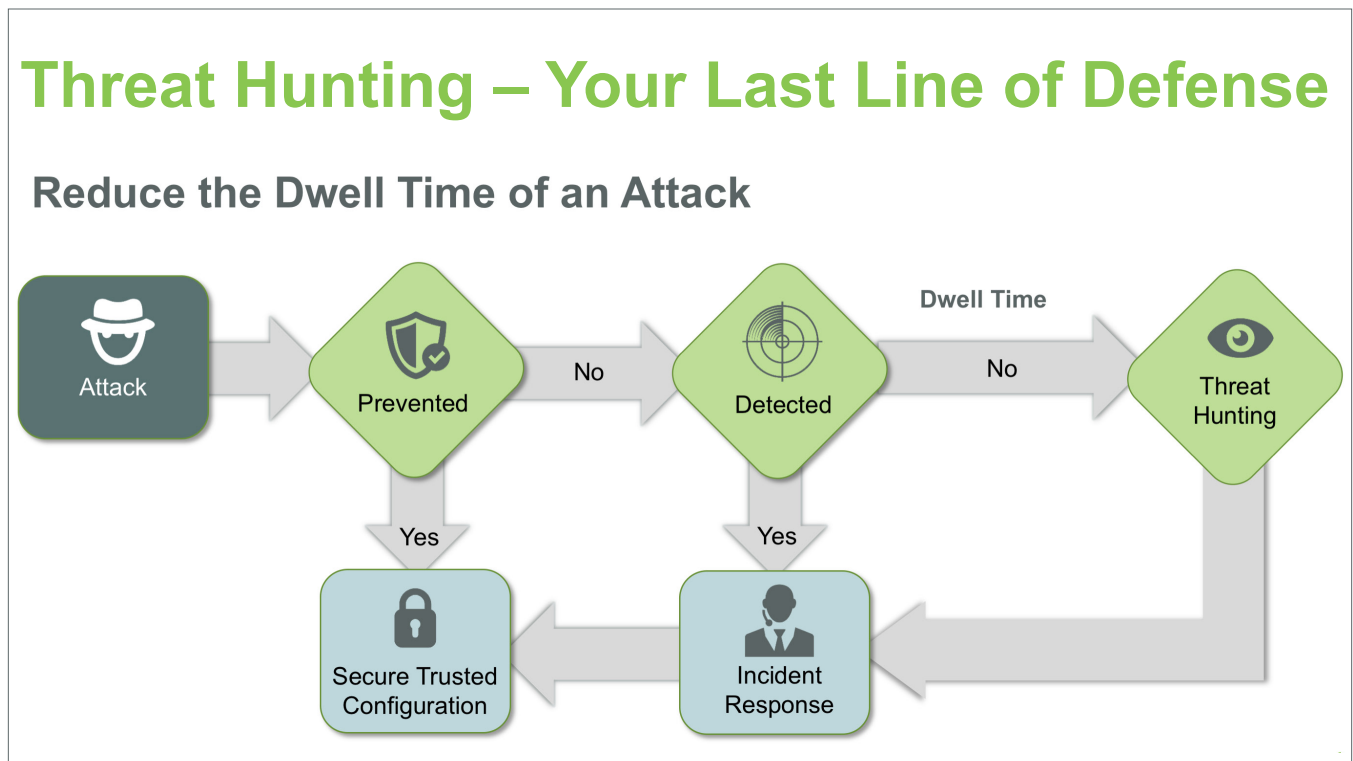
While an incident responder goes into action when notified of a potential attack, a threat hunter flips that model on its head to proactively look for trouble and indicators of compromise before an alert occurs. Like the best criminal investigators, a threat hunter has to have threat intelligence about techniques, tactics, and procedures (TTPs) of the adversary to start the hunting process. Additionally, organizations need to support threat hunting with the right tools and data. Rich metadata collected from network sensors, endpoints, and cloud environments allows for cross-session analysis and multi-faceted and behavior analysis; and is critical for post-breach investigations of the unknown.

Companies are advised to have mature threat detection and alert triage skills before considering threat hunting. Hypotheses for threat hunting mainly come from the expert intuition of security analysts, and include past incident data, threat actor capabilities, red team and threat simulation artifacts, threat intelligence, and anomalies. For example, traffic to dynamic DNS sites indicating exfiltration or command and control, known

file names with uncommon paths masquerading as Windows processes, or the presence of RAR files for attack staging.

The most popular tools for threat hunting are endpoint detection and response (EDR) and network traffic analysis (NTA) with automation capabilities to help gain direct visibility of unknown advanced attacks. Hidden behind these tools is metadata for real-time and retrospective analysis by content and context not visible within restricted NetFlow data, logs, and events.

With proactive threat hunting, organizations can shift from being on their heels, waiting for the inevitable and hoping to detect it before real damage is done, to being on their toes. But, with an overall cybersecurity skills shortage, don't push your SOC team to threat hunt while they are still battening down the hatches and ensuring basic security protocols are met. If you need to get on the hunt, and don't have the manpower, consider an MDR service with access to the metadata data needed to be 'on the hunt.'



4 Engage Managed Detection and Response Services

The cybersecurity skills gap continues to widen and CISOs are feeling the pressure. In its 2018 CIO Agenda Survey, for example, Gartner reports that only 65% of CISOs say they have a cybersecurity expert. Capitalizing on the situation, attackers are increasingly striking outside of normal business hours. Threat actors, for example, may execute financial attacks on Saturday morning knowing security resources are out of office or lightly staffed. While off-hours attacks are driving the need for round-the-clock coverage, 24/7 coverage requires eight to twelve FTEs – a requirement that many security operations teams simply cannot staff.

Finding, hiring, and retaining talent are distinct challenges, exacerbated by the fact that the skill sets are highly specialized. Most security teams have skills matching preventive security defense technologies. Skills that align to post-breach detection and response are needed and are in short supply. Such skills often develop from years of incident response and penetration testing or red team versus blue team exercises and threat hunting. The lack of security talent and resources on staff is driving a demand for Managed Detection and Response (MDR) services. While managed detection and response may sound similar to the Managed Security Services (MSS) that organizations have turned to in the past, MDR is more than a name change.

Gartner reports that **only 65% of CISOs say they have a cybersecurity expert.**

Managed Security Services Providers (MSSP) have ignored customer requests for services beyond compliance and security monitoring, enabling the rise of services focused on detection and response. Managed detection and response is a service that augments an organizations' existing security infrastructure with a third-party team of cybersecurity specialists who ensure an organization has 24/7 threat detection and response coverage. Managed detection and response improves advanced threat detection and incident response capabilities via a turnkey approach using a curated technology stack that often unifies network visibility, endpoint detection and response, and deception to detect threats that have bypassed preventive defenses and security controls. Many MSSPs are calling themselves MDRs, but while they provide basic security capabilities, they lack the necessary incident response expertise and in-house technology to detect and respond to advanced threats.



While off-hours attacks are driving the need for round-the-clock coverage, **24/7 coverage requires eight to twelve FTEs** – a requirement that many security operations teams simply cannot staff.

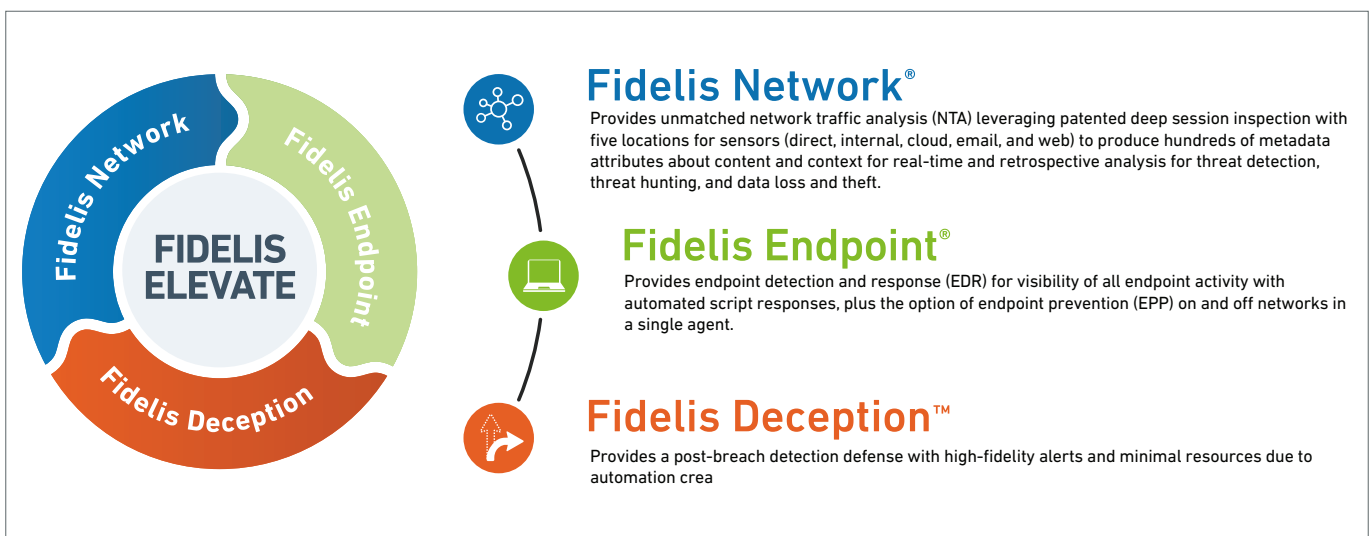
Fidelis Elevate™ – Your Force Multiplier for Automating Detection and Response

Fidelis Elevate™ provides a rich source of metadata for real-time and retrospective analysis by integrating network visibility, network data loss prevention, deception, and endpoint detection and response into one unified solution to deliver automated threat detection and response across networks, endpoints, cloud and enterprise IoT environments. It provides threat visibility and intelligence with content and context to help organizations quickly address cyberattacks across the entire threat lifecycle – from initial intrusion to exploitation to data theft – as well as hunt for unknown threats deep within the network.

Detected threats are presented as a conclusion that was determined by validation from network to endpoint, contextual enrichment, and correlated threat activity to enable the analyst to take rapid, responsive and often automated action. By delivering comprehensive visibility, alert validation, and increased speed to respond, it enables security teams to focus on the most urgent threats and protect sensitive data rather than spending time investigating and triaging thousands of alerts.

Fidelis Elevate includes the following security products, which can be deployed and activated as a complete Elevate platform or individually based on an enterprise's needs:

- **Fidelis Network:** Provides unmatched network traffic analysis (NTA) leveraging patented deep session inspection with five locations for sensors (direct, internal, cloud, email, and web) to produce hundreds of metadata attributes about content and context for real-time and retrospective analysis for threat detection, threat hunting, and data loss and theft.
- **Fidelis Endpoint:** Provides endpoint detection and response (EDR) for visibility of all endpoint activity with automated script responses, plus the option of endpoint prevention (EPP) on and off networks in a single agent. Real-time detection uses behavioral rules and indicators while security analysts can also use third-party feeds and custom rules for threat detection, as well as hunt for threats directly on the endpoint, in both the file system and memory, using YARA and Open IOC.
- **Fidelis Deception:** Provides a post-breach detection defense with high-fidelity alerts and minimal resources due to automation creating, adapting, and freshening deception layers. Auto-generated decoys and breadcrumbs make deception deterministic, luring attackers from compromised hosts to decoys. Deception also provides detection for legacy systems, shadow IT, and enterprise IoT devices.



Fidelis Network. Fidelis Network automates threat detection and response while preventing data leakage across the network, cloud, email, and web. It bi-directionally scans all network traffic to reveal network and application protocols, files, and content; and automatically decodes and analyzes traffic to detect advanced threats and unauthorized data transfers. Patented Deep Session Inspection (DSI) as well as Deep Packet Inspection (DPI) give unique visibility of deeply embedded content and context across all ports and protocols. It also captures and stores over 300 attributes of standard metadata, plus enhanced metadata, such as custom tags, to provide rich information for automated and manual threat detection and threat hunting with up to 360 days of retrospective analysis. More information on the power of metadata can be found in the white paper: [What's Hiding Within Your Metadata? How to Decode Your Network's Deepest and Darkest Secrets.](#)

Fidelis Network solves the 'My analysts are not experienced enough' problem by providing assisted and automated investigation. Detected advanced threats are presented as a conclusion determined by automated validation, contextual enrichment, and correlated threat activity. Validated, enriched, and prioritized alerts dispel uncertainty, enabling the analyst to take rapid responsive and automated actions rather than spending time to gather evidence from a variety of security systems and sources.

Fidelis Endpoint. Fidelis Endpoint provides visibility into all endpoint activity, including process actions, logged in users, registry writes, file system activity, and memory. It monitors endpoints in real-time and retrospectively, on and off the network, and records key events, allowing security teams to see a timeline of suspected incidents. Fidelis Endpoint drives real-time detection through the use of behavioral rules and indicators provided by Fidelis Insight threat intelligence. Security teams can also use third-party feeds and custom rules for threat detection, as well as hunt for threats directly on the endpoint, in both the file system and memory, using YARA and Open IOC.

Fidelis Endpoint equips security teams to confidently detect, respond to, and resolve security incidents in a fraction of the time it takes using traditional approaches. Security analysts are also able to respond to endpoint activity faster by integrating with SIEMs, NGFWs, alerting tools, and more and by accessing a large library of pre-defined as well as custom response scripts. Fidelis

Endpoint further aids security teams by automatically kicking off response workflows and scripts for automated remediation or deep analysis actions when suspicious activity is detected.

Fidelis Deception. Fidelis Deception significantly reduces dwell time by providing a low-risk, low-friction internal alarm system to detect post-breach attacks and malicious insiders. Deception defenses provide a proactive opportunity to lure, detect, and defend early within post-breach compromise incidents with no risk to resources or data, or impact to users and operations. Fidelis Deception automatically discovers the environment and auto-generates decoys that have profiles, services and activity matching the environment for active deception layers. It deploys decoys of key assets, services, and fake data, and then makes deception deterministic by setting up breadcrumbs on real systems likely to be compromised, thus leading attackers to decoys. High-fidelity alerts come from decoys, breadcrumbs, AD credentials, MITM, and poisoned data with network traffic analysis and telemetry data for investigations. Lastly, Fidelis Deception automatically adapts the deception environment to network changes as they occur to remain synchronized for assets, resources, and services. For a deeper discussion on deception and how to apply deception mechanisms to detect sophisticated cyber attacks, download the white paper [Capture the Flag with Deception Defenses.](#)

Fidelis Managed Detection and Response Service.

Fidelis Managed Detection and Response can augment existing SOC operations or act as a trusted security team running the Fidelis Elevate curated security stack and leveraging the DNA of Fidelis' rich metadata. Fidelis MDR delivers industry-best talent and solutions to proactively hunt for threats, fully investigate, respond to detected threats, and stop attacks and data theft. The 24x7 Threat Analysis Center is staffed by a team of highly trained security analysts and incident responders who have supported over 4,000 incident response cases and provided expert testimony in more than 100 court cases for commercial and federal clients. Fidelis MDR includes deception to lure adversaries away from critical assets and sensitive data to decoy environments that look and feel real. By ensuring faster threat detection and response to advanced threats and preventing data theft, Fidelis MDR enables organizations to focus on their core business.

Summary

Many organizations are discovering, too late, that despite the prevalence of prevention techniques and tools at their disposal, attacks are succeeding in growing numbers. The inability to see, detect, and respond to modern, complex, post-breach attacks hampers the effectiveness and efficiency of security operations teams to reduce dwell time. That's why Gartner has recommended organizations shift from a prevention-focused defense to one that prioritizes detection and response.

Curated security technology stacks designed for detection and response using a rich metadata model with visibility across network, cloud, and endpoint to understand content and context is the leading analyst recommendation for advanced threat detection and response. Combining deep and broad visibility on both the network and the endpoint with fast, comprehensive detection with specific validation and mature response and prevention capabilities can enable security operations teams to employ detection techniques such as network traffic analysis, endpoint forensics, and payload analysis, as well as combine techniques. By leveraging automation for detection, investigation, and response processes, as well as automating deception deployments, organizations can improve the performance of the security team while effectively responding to the cybersecurity skills drought. Finally, security leaders are encouraged to not go it alone. The engagement of managed detection and response services can augment an existing SOC with proactive threat hunting or act as a trusted security team.



Fidelis Cybersecurity is a leading provider of threat detection, hunting and response solutions. Fidelis combats the full spectrum of cyber-crime, data theft and espionage by providing full visibility across hybrid cloud / on-prem environments, automating threat and data theft detection, empowering threat hunting and optimizing incident response with context, speed and accuracy.

By integrating bi-directional network traffic analysis across your cloud and internal networks with email, web, endpoint detection and response, and automated deception technology, the Fidelis Elevate™ platform captures rich metadata and content that enables real-time and retrospective analysis, giving security teams the platform to effectively hunt for threats in their environment. Fidelis solutions are delivered as standalone products, an integrated platform, or as a 24x7 Managed Detection and Response service that augments existing security operations and incident response capabilities. Fidelis is trusted by Global 1000s and Governments as their last line of defense. Get in the hunt. For more information go to www.fidelissecurity.com.