

White Paper

HardenStance

What to Expect from MDR & MDR Providers

By Patrick Donegan, Principal Analyst, HardenStance

Sponsored by

AT&T Cybersecurity  Eleven Paths  FORTINET.  Secureworks®

June 2020



HardenStance

*"Trusted Research, Analysis and Insight in IT
& Telecom Security"*

Executive Summary

- Greater sophistication of threats, alert overload, and a shortage of top talent to run security operations, make detection and response increasingly challenging.
- Failing to detect, contain and respond to cyber threats is a high risk for organizations. Alert triage, investigations and response require a lot of expertise.
- The case for Managed Detection and Response (MDR) can be very compelling. There are a variety of MDR providers but MDR is always a service, not a product.
- For an MDR provider to meet a customer's cyber security goals, expectations regarding each party's responsibilities must be tightly aligned. Buyers must hold their own feet to the fire as well as those of their MDR providers.

In a Changing World, Basic Tenets Still Hold

For CISO teams evolving their security posture in lock step with the demands of digital transformation while facing accelerating cyber threats, life was hard enough before COVID-19. The pandemic has heaped on new risk from greater, faster, changes in remote enterprise networking. It has also made employees more anxious, rendering them more vulnerable to risky on-line behaviour.


However tough cyber security has been in the past, it's a safe bet that it's not going to get any easier. Faced with these challenges, CISO teams are having to sustain investment in cyber security at a time when in many cases their organization's revenues have slumped. In some cases, business revenues have collapsed.

Meantime, cyber adversaries aren't just largely immune from the travails of mainstream legitimate business activity. To the contrary, attackers can even feel emboldened by the new 'opportunities' that this crisis, and its lasting legacy, is creating for them.

Looking out at these challenges, the following basic tenets of cyber security remain true:

- **Security is only as strong as the weakest link.** That's the reason it's not possible to elevate the importance of any one of the five cyber security domains depicted in **Figure 1** above any of the others. In the context of this paper, weak identification and protection make improving detection and response harder, if not impossible.
- **While the big picture is always critical, a complex discipline like cyber security also needs breaking down into more manageable components.** This allows organizations to apply the right focus to execute on each component part with maximum efficacy and efficiency.

Figure 1: Detection & Response in the NIST* Cyber Security Framework



Source: HardenStance

*National Institute of Standards and Technology

Weak identification and protection layers make detection and response a lot harder.

This White Paper focuses on the 'Detection' and 'Response' components of enterprise cyber security. It examines the state of the art in threat detection and aligning organizational processes to trigger rapid responses to minimize the damage that cyber threats can cause. In particular, this White Paper explores the value proposition of Managed Detection and Response (MDR) service providers and how enterprises can get the most from partnering an MDR provider.

The Rise of the 'MDR' Category in Cyber Security

Focused on managing perimeter devices like firewalls, and operated either by an organization itself or a Managed Security Service Provider (MSSP), detection and response existed long before NIST's cyber security framework was released in 2013. But MDR as a discrete solution category didn't exist until Gartner defined it in 2016.

Before looking at how the definitions and scope of 'detection' and 'response' have evolved, it's important to consider some of the underlying drivers for the creation of MDR as a new category. The first of these is changes in the threat landscape:


- A subset of threats have become more sophisticated and now easily evade traditional perimeter controls to target valuable and sensitive data.
- Threats manifest themselves in East-West traffic within an organization's environment and North-South between the organization and the global Internet.
- Nowadays, some advanced threats don't move sequentially from left to right along the attack 'kill chain' either. As shown in **Figure 2**, some will do things like carry out an execution to steal data and jump ahead into an exfiltration phase before returning to privilege escalation to set up a long-term persistence mechanism. These Advanced Persistent Threats (APTs) secure multiple footholds to retain a presence even after a target organization thinks it has successfully evicted them.

Digital Transformation has Exposed Organizations to Greater Risk

Organizations have also exposed themselves to greater risk via their own digital transformations. Cloud migration, Bring Your Own Device (BYOD), enterprise mobility, remote networking and the Internet of Things (IoT) all increase the volume and variety of devices that can access enterprise data and applications – as well as the variety of environments in which these data and application assets must be hosted and protected.

On the whole, the response to these new risks from industry at large, governments and the cyber security industry itself, has been flawed. The cyber security industry has focused too much investment on adding to the already-huge variety of cyber security products and product categories, many of which overlap and compete with one another.

Figure 2: Linear & Non Linear Attack Progression along the Kill Chain



Source: HardenStance

Nowadays some advanced threats don't even move sequentially along the attack 'kill chain'.

Not enough has been done to fill the shortfall in demand for cyber security professionals, leading to fierce competition for talent. Hence the core problem of cyber security today is that security analysts are overloaded with alerts from often dozens of different tools. They can't quickly detect the most significant threats from the huge volumes of alerts that pose little or no risk at all. All too often they either respond late to high risk threats, by which time the attack has already been partially successful. Or worse they can't respond at all, resulting in large-scale damage. According to the Ponemon Institute's 2019 'Cost of a Data Breach' survey, the average time to identify a breach in 2019 was 206 days. The average time to contain one was 73 days, for a total of 279 days.

MDR assumes a subset of threats will inevitably penetrate perimeter controls. Hence it assumes there are only two types of organizations – those that have been breached and those that don't know they've been breached. The goal of an MDR provider is to harden an organization's security posture and in particular to lower its Mean Time To Detection (MTTD) and Mean Time To Respond (MTTR) in a Service Level Agreement (SLA). These are addressed in more detail in the concluding section of this paper but at a high level, the shorter the MTTD, the harder it is for attackers to cause an incident. Likewise, the shorter the MTTR, the less harm an incident can cause.

The goal of an MDR provider is to improve an organization's security posture and in particular to lower its MTTD and MTTR.


The Origins of MDR can be found in EDR

A key turning point in the origins of MDR was the successful adoption of the Endpoint Detection and Response (EDR) product category. Keeping in mind that an endpoint – a workstation, desktop PC, laptop, mobile phone, tablet, server or IoT 'thing' - is nearly always the end target of a cyber attack, EDR transformed detection and response disciplines by generating far more granular visibility, correlation and contextualization around endpoint activity than any OS or Antivirus (AV) logs had ever been capable of.

In modern-day security operations, the curation of EDR information and correlation with other information sources is key to making informed choices of which threats to investigate and how to investigate them. In that sense, efficiencies in detection and response rates enabled by EDR products are a baseline requirement for MDR. This White Paper makes the case for layering on a lot more capability on top of this for a comprehensive MDR service as offered either by an MSSP with its own MDR line of business or a start-up MDR provider. Where an EDR product vendor manages their product as a managed service integrated within a customer's own Security Operations Centre (SOC), that can and should be considered a baseline MDR service as well.

As shown in **Figure 3** below, there is a lot of diversity in today's MDR market in terms of both supply (different categories of MDR provider with very different backgrounds) and demand (different types and size of companies with very different risk profiles, vulnerabilities to particular threat actors, and hence requirements for MDR providers).

Figure 3: Users Need to Map their Requirements to the Right MDR Provider



Source: HardenStance

Two further points need making with respect to the definition of MDR:

- As often happens with a hot technology acronym, a big challenge with identifying the right MDR service is that some vendors are adjusting, expanding or even distorting the 'MDR' term to ensure their own products get bathed in some of its marketing sunshine - even though they don't deliver an end to end MDR service.
- The term 'XDR' has started to be used by some product vendors, often in conjunction with MDR. For example, different vendors have used 'XDR' branding to categorize Network Traffic Analysis (NTA), Next Generation Firewall (NGFW), Security Orchestration Automation and Response (SOAR) and EDR products, either as stand-alone products or as part of an 'XDR' suite. The use of XDR can either be wholly complementary or adjacent to MDR. So an EDR product forming part of an XDR suite can form part of an MDR service. XDR-branded NTA products can also be part of an MDR service.

Good detection and response cannot be implemented without good identification of risk and protective measures.

A Guide to Comprehensive Detection & Response

As summarised in **Figure 4**, this section sets out the requirements for a first class Security Operations Centre (SOC) dedicated to round the clock monitoring, detection and response. Given the round the clock requirement, a team of six to eight analysts tends to be a bare minimum requirement.

The following are generic requirements independent of whether a SOC is built entirely by an organization itself; wholly outsourced to an MDR service provider; or built as a hybrid solution combining an organization's own building blocks with those of an MDR partner. Where relevant, reference is made to the specific challenges MDR providers face in integrating into a customer's environment and workflows.


As a foundation for state of the art detection and response, organizations must correctly assess, and regularly review, what their most critical data assets are and the risks to them. Organizations must also implement the best possible protective mechanisms like employee awareness training and access controls.

Weakness in either of these foundational pillars of the NIST framework (see **Figure 1**) introduces weaknesses into detection and response. Protecting the wrong things, or protecting the right things in the wrong way, makes it much harder to quickly detect and respond to high-risk threats when they get through. Good detection and response cannot be implemented without good identification of risk and protective measures.

Comprehensive Visibility and Continuous Monitoring

The first requirement is to have complete visibility of activity across an organization's IT estate, including insight into all its endpoint, network, cloud and SaaS assets. 'Shadow IT' must also be included, since any one dark corner that can't be seen into poses a risk. EDR has transformed visibility into endpoint activity but it doesn't provide any insight

Figure 4: State of the Art Detection & Response



Source: HardenStance

Adjusting to the ambiguous outcomes generated by AI algorithms that are increasingly embedded in security tools is proving to be a steep learning curve.

into many connected IoT 'things' like printers, Wi-Fi Access points or an employee's Fitbit because they can't support EDR endpoint agents. Moreover, endpoints are still just one component in a state of the art approach. Traditional network telemetry and events generated by switches, routers, Next Generation Firewalls (NGFW) and Intrusion Detection Systems and Intrusion Protection Systems (IDS/IPS), as well as data or application security controls, are key inputs.

Visibility and detection needs extending to business critical assets as they migrate to the cloud as well as to SaaS applications that are common attack targets. The impact of COVID-19 is driving still greater focus on remote networking assets like home routers. This patchwork of environments, each with different security models, needs continuous security monitoring 24/7, 365 days a year.

Data Ingestion, Alert Triage and Immediate, Automated Response

Whilst it is absolutely critical, the challenge with comprehensive visibility and continuous monitoring is that it heaps still greater volumes of events onto the security team. These feed into some kind of aggregation point or single pane of glass which has traditionally been a Security Incident & Event Management (SIEM) platform.

It's from this point on, starting with tuning aggregation platforms to minimize false positives, and the way events are triaged, that detection and response disciplines are undergoing a transformation – and MDR is playing a critical role in driving that.

In many cases today, the only way Level 1 security analysts can cope with the volume is simply by ignoring large swathes of the alerts they see, which is a substantial risk in itself. Achieving stringent MTTD and MTTR targets mandates handling all alerts. It requires identifying the vast majority as 'known', low level, threats and mitigating them through machine automation. That frees up Level 1 analysts to focus on the subset of alerts that remain; mitigate or remediate those they can; and escalate those that look like they could represent a significant risk to Level 2 Incident Response (IR) analysts.

Across the incident triage, investigation, and Incident Response (IR) disciplines, analysts need to be able to use the latest detection and response technologies as well as traditional tools like SIEMs. EDR has already been flagged as critical. As mentioned, another key new technology is SOAR, which pulls data from multiple different security platforms and automates IR workflows.

Capabilities that look at attacks from the perspective of specific users such as User and Entity Behaviour Analytics (UEBA) are being layered on here. Deception technology can also be very useful. Machine Learning (ML) and Artificial Intelligence (AI) also need to be exploited for computational tasks that would take an analyst hours or days. As shown on page 7, strategic, operational and tactical threat intelligence must be gathered, interpreted and applied throughout this environment – in real time or near real time.

Many of the Newer Detection Tools are much Harder to Manage

Many of these new technologies are a lot more challenging to tune than traditional perimeter controls. Whereas the latter tend to be used 'conservatively' – erring on the side of letting things through to keep the business going – EDR is inherently more invasive. In applying it to unearth hidden threats, the risk with EDR is that if it isn't really well tuned it will either generate many more false positives or miss an embedded threat altogether. Adjusting to the ambiguous outcomes generated by AI algorithms that are increasingly used to enhance security tools is proving to be a steep learning curve.

Rather than try building their own state of the art SOC, organizations can leverage the SOC investments that MSSPs and MDR start-ups have made to support these requirements. Their platforms integrate third party vendors and proprietary components, though they should provide choice in key product categories like EDR.

Differentiation in Strategic and Operational Threat Intelligence

Accessing threat intelligence expertise should be a key driver for using an MDR provider. The breadth and scale of what MDR providers see, and have experience of applying, while monitoring hundreds of customer environments across endpoint, network, cloud and SaaS assets, is way beyond what most organizations can see themselves.

Not all threat intel is created equal, though. Most technical intelligence – Indicators of Compromise (IoCs) like bad IP addresses and URLs – is either open source or freely commercially available (to both defenders and attackers). It quickly loses its value, hence it's only basic table stakes for an MDR provider.

Isolated technical intelligence artefacts have limited value. Enriching them with other contextual inputs in the context of the organization's specific risk profile is what adds value. MDR providers can differentiate in the way they generate, correlate, and apply operational and strategic intelligence and the ease with which analysts can access it from their platform. Operational Intelligence centers around the 'how?', 'where?' and 'when?' of cyber security incidents and has to be augmented by human expertise. All the major MDR providers now leverage the detailed mapping of the Tactics, Techniques and Procedures (TTPs) used by major threat actors in the MITRE ATT&CK Framework.

MDR providers can further differentiate by means of strategic intelligence which focuses on the 'who?' and 'why?'. This tracks and analyzes changes in funding, motivations and behaviours of major threat groups to achieve still greater intimacy. It tracks their activity in the real physical world as well as in the cyber realm and watches for changes in one domain impacting the other. The best MDR providers invest a lot in strategic intelligence, working closely with international law enforcement agencies and engaging in activities such as infiltrating these threat groups in order to understand their workings from the inside.

The best MDR platforms enable the provider's analysts to pivot easily and efficiently between different tasks in an investigation. They give a customer's team access to the platform and allow communications with their own team in real time. They also provide customers with a visual record of all the investigations the MDR provider has conducted and how long it took to complete each one.

An MDR partner should help source freely available defensive playbooks and, if required, support customers in customizing them to their unique environment.

Initial Response and Incident Response (IR)

A cyber security incident is a breach of a security policy. Response contains and neutralizes the impact of incidents. As shown in **Figure 1**, Recovery is a distinct discipline. Some MSSPs that are MDR providers play in that segment but pure-play MDR providers typically don't.

From the perspective of real-time incident management, any response process begins with an initial detection. Way before any detection is ever made, however, state of the art Incident Response begins with meticulous preparation, documentation and rehearsal of defensive playbooks. These map out an organization's plans for responding to generic or commodity threats, as well as the high risk threats identified in its detailed risk assessment. An MDR partner should help source freely available defensive playbooks. If needed they can help organizations customize them to their unique environment as well.

On its own, or assisted by an MDR partner, a customer organization's security team can respond to many incidents by themselves. High risk incidents often require coordination across security, IT and other departments such as HR, Finance and Legal. Creating and maintaining defensive playbooks is key to preventing an organization's understanding of its IR processes from being closely guarded as 'tribal knowledge' by a select few – then being lost when those individuals leave the organization.

One of the key areas where MDR providers best demonstrate their value is in identifying the impact of an incident to date – for example how localised it is, how many vectors are under attack. That then serves as the basis for determining next steps – whether a

containment action should be initiated, whether an investigation should be started and whether threat hunting steps should be initiated.

There are four types of Incident Response that MDR providers can serve up. These are:

- 1 **Remotely inform:** notify and leave the customer to determine the response.
- 2 **Remote response recommendation:** notify the customer and recommend action.
- 3 **Remote response execution:** implement a response on behalf of the customer.
- 4 **On-site boots on the ground:** execute a response and remediation at the customer's premises.

Four things should be noted about these different response types:

- The second and third response types, integrated into a customer's own workflows, are generally preferable. Certainly, a risk with a poor choice of MDR provider is that too many responses will consist of 'throwing' alerts 'over the wall' without context or recommendation, leaving the customer to figure out what to do. That said, there are many examples when that type of notification is the right one. For example, a user policy violation where only the organization itself can determine whether an individual's permissions should be changed. Similarly, it is not possible for MDR providers to execute a response themselves where that would require full admin control of an Active Directory that third parties are usually not given access to.
- As an end to end definition, the term 'Incident Response' covers everything from preparation to initial detection, alert triage and analysis, containment and neutralization. However, the term has tended to be closely associated with a distinct service category focused on the 'boots on the ground' response where high severity incidents are dealt with on a customer's premises. Many IR specialists only serve this segment. The boundaries between remote and boots on the ground IR are becoming more porous with the move to the cloud and greater remote networking arising from the COVID-19 pandemic. But the distinction still tends to be a core feature of the way IR is delivered as part of an MDR service.
- Many MDR providers provide all four types of response and there are important advantages of consistency in a single provider managing the entire process and closing the loop for the customer. That said, some MDR providers only offer initial remote responses that are sometimes referred to as 'lower case r'. They leave customers to choose a different partner for boots on the ground IR.
- MDR providers typically offer a basic quota of boots on the ground IR hours as part of a core service. Even if they want to take the risk of not investing in these up front in their contract – or choosing another focused IR provider if or when the need does arise – organizations should at least look for zero cost or zero commitment contracts with their MDR provider. These ensure legal and financial aspects of a major IR engagement are agreed before a high severity incident occurs. That avoids losing days of valuable time putting an IR deal together when time is of the essence.

The boundaries between remote and boots on the ground IR are becoming more porous with the move to the cloud and greater remote networking arising from the COVID-19 pandemic.

Threat Hunting Driven by Threat Intelligence Expertise

A core goal of better detection and response is to enable organisations to free themselves from merely reacting to threats and incidents and get on the front foot by proactively seeking out threats in their network. Threat Hunting has come to feature front and center in this and is becoming an increasingly complex and sophisticated discipline. Any leading MDR provider should excel at it.

Threat hunting can take two forms. A manual hunt looks for the presence of Indicators of Compromise (IoCs) against the very latest Tactics Techniques and Procedures (TTPs). Analysts can take multiple approaches in their manual hunts, from posing a hypothesis and testing it in the environment to using the tools of forensics and Incident

Response to pen testing. Automated hunts can also be triggered at regular intervals in pursuit of well-known threats, albeit since they have to be led by experienced and skilled SOC analysts, they cannot be fully automated.

As well as identifying and evicting hidden threats in an organisation's environment, threat hunting can also yield new IOCs that can be fed back into an organisation's monitoring environment. While all types of threat intelligence are widely used throughout detection and response, successful threat hunting is perhaps more dependent on expertise in applying it than any other sub-discipline within cyber security.

Headline Factors Can Seem Alluring

Here are the main reasons outsourcing some or all of their detection and response capabilities to an MDR provider is something most organizations should be considering:

- avoid high annual capex on cyber security technology;
- avoid high cost of managing vendor relationships and the vendor integration effort;
- avoid having to master advanced cyber operations and threat intelligence;
- avoid fierce competition to hire, motivate and retain top talent to run the operation.

Many large organizations are successful doing these things themselves – but at a high cost. For example, big banks spend hundreds of millions of dollars a year each on cybersecurity. Based on discussions with MDR providers, and assuming a North America-focused business with a thousand users, HardenStance research points to a ballpark price for a three-year MDR service agreement in the range of around \$450,000 - \$750,000. That's around \$150,000 - \$250,000 a year, depending on the specific service mix and the specific MDR provider.

It's easy for executives to review the high level factors cited above, find them persuasive, and feel at least somewhat favorable towards engaging an MDR partner. Big outsourcing decisions need a lot more detailed consideration, though. In particular, organizations need to consider exactly how they will hold an MDR provider's 'feet to the fire' to improve their security outcomes. Just as importantly, they need to ask themselves how willing they are – what long term commitments they themselves are willing to make – to hold their own feet to the fire to achieve those same ends.

Holding an MDR Provider's 'Feet to the Fire'

There are hard and soft ways to hold an MDR provider's feet to the fire. Among hard metrics, the most important are SLAs that specify agreed Mean Time To Detect (MTTD) and Mean Time To Respond (MTTR) metrics.

There are several stages in an investigation and a response. These include the time it takes the MDR provider to make an initial detection; time taken to report that detection to the customer (with or without context); the time it takes the customer to view the report; and the time it takes the customer to consider the investigation complete.

An example of one MSSP's current SLA for 'Time To Report an Alert to a Customer Following a Severity Assignment' is shown in **Figure 5** overleaf. Buyers should recognize, however, that there aren't any industry standard approaches to sub-dividing and measuring each individual 'Time To' component of MTTD and MTTR. Hence mapping an organization's unique requirements to the different SLAs offered by different MDR providers isn't an apples-to-apples comparison. It requires a lot of detailed attention.

SLA commitments to MTTR are more customized because they are so dependent on the level of integration that organizations allow the MDR provider to have with their own environment. These SLAs reflect the extent to which a customer is willing to allow the MDR provider to reach into its own IT environment and implement configuration changes on the customer's behalf. Buyers should demand a commitment to continuous

SLA commitments to MTTR are necessarily more customized because they are so dependent on the level of integration that organizations allow the MDR provider to have with their own environment.

Figure 5: An MDR SLA for Time to Report an Alert to a Customer Following a Severity Assignment

Alert Severity	Description	Initial reporting timeframe following severity assignment
Critical	An individual gains logical or physical access without permission to a customer’s network, system, application, data, or other resource.	5 minutes
High	An attack prevents or impairs the normal authorized functionality of networks, systems or applications by exhausting resources. This activity includes being the victim or participating in the DoS.	15 minutes
Medium	Successful installation of malware (e.g., virus, worm, Trojan horse) that infects an OS or application. Customers are not required to report malicious logic successfully quarantined by AV software.	60 minutes
Low	A person violates acceptable computing use policies.	60 minutes
Low	An activity that seeks to access or identify a customer’s computer, open ports, protocols, service, or any combination for later exploit. This activity does not directly result in a compromise or DoS.	60 minutes

Source: Current MDR provider (anonymized at the provider’s request)

improvement. For example, is the MDR provider willing to commit to improving a specific SLA commitment by, say, 20% in year three? Within the MDR provider’s platform, customers need to be able to view the provider’s performance in each investigation against the target SLA as well as track the average performance against the target SLA.

Buyers shouldn’t focus exclusively on hard metrics. At both the time of vendor selection, and on an ongoing basis, they should also pay close attention to factors whose impact might not be directly measurable on a weekly or monthly basis but do nevertheless impact the MDR provider’s performance and their organization’s cyber security posture.

There are a number of principles to be guided by here, including the following:

- **Investigate a provider’s human resources policy** as well as its technology roadmap. Stress levels and staff turnover rates among SOC analysts are notoriously high. How an MDR provider supports the mental health of its staff as well as working conditions such as the length of shifts can be as important as an employee’s salary and benefit packages.
- **Focus on learnings from investigations that have been performed into high risk threats.** What has been learned and what does the MDR provider see as implications for where its roadmap and SLA commitments should go next?
- **Be wary of MDR providers over-promising interaction with high cost analysts,** seemingly independent of the issue at hand, and also be wary of providers that don’t center their value proposition around the higher margin end of the business where complex investigations are needed and where responses are orchestrated and automated. Historically, managed security services have had difficulty turning a profit. Overcommitting on expensive human resources and under-committing to more challenging, but higher margin, market segments are two of the key underlying reasons for that.
- **As with anything in cyber security, trust but verify the MDR provider.** Organizations should run regular red team exercises against their own systems in order to verify for themselves exactly how well the MDR provider’s processes are working. Since hacking into an MDR provider can be a way into its customers networks, organizations should also require verification of how the MDR provider secures its own environment.

As with anything else in cyber security, trust but verify the MDR provider.

Buyers Must also Hold their own 'Feet to the Fire'

Wanting to hold an MDR provider's feet to the fire comes naturally. What's much harder – but just as important – is for organizations to be entirely honest with themselves about what is needed on their side to achieve the desired security outcomes – and whether they themselves are truly committed to doing what's required. Whether the organization is or not depends on the answer to one question: how serious is it about moving forward with digital transformation while also reducing the risk posed by cyber threats?

The importance of not throwing problems 'over the wall' applies in both directions. Interested buyers must first recognise that hiring an MDR provider expecting to save time as well as money will not improve security outcomes. Cyber security is a team sport. An MDR provider's ability to achieve its customer's goals are heavily influenced by the quality of day to day operational alignment between the processes of the MDR provider and the customer organization's security, IT and other teams. Engaging an MDR partner effectively could well require the customer organization as a whole to invest as much, if not more, of its own time on cyber security than it ever did previously.

Organizations Must Commit to Aligning with their MDR Provider

This remains critical throughout the duration of a multi-year MDR contract, not just at the outset. Organizations preparing to partner an MDR provider must recognise the obligations on their part to be willing to align their organization's processes to the MDR provider's as well as expecting the MDR provider to align with theirs.


The most obvious manifestation of this is that outsourcing large parts of capex, integration and vendor relationship management costs – and gaining access to a state of the art MDR platform – requires writing off some of one's own investments. An MDR provider should be expected to integrate with some, but by no means all, elements of a customer organization's existing security infrastructure.

Organizations should also abide by the following example principles in enabling their MDR provider to meet their contractual obligations:

- Give the MDR provider ongoing intimacy with, and a feedback loop into, the definition and protection domains of the organization's security framework as well as forewarning of upcoming IT projects. The MDR provider is not accountable for a data breach if the root cause analysis shows it didn't have visibility of a compromised endpoint and therefore hadn't instrumented it. Upcoming monitoring requirements should be factored in at the time of agreeing a contract – not a year into it.
- Allow the MDR provider to lead in jointly re-writing many aspects of detection and response procedures spanning the two organizations. The MDR provider's expertise is derived from hundreds of customer experiences. An internal security team's legacy processes will not map one to one to the MDR provider's new ones. The internal team should understand at the outset which of its traditional processes are redundant, why, and what the MDR provider is replacing them with. This will reduce potentially corrosive disputes in the future over the importance of alerts that were previously considered significant but are now being disregarded.
- Define clearly how the new security policies and requirements agreed with the MDR provider will be implemented internally within the organization, according to specific milestones. Either the security agenda needs to be formally prioritized by the organization's IT leadership or its security team needs to be required and enabled to execute some of its agenda within the IT domain. If this delivery model isn't clearly defined internally, the roadmap agreed with the MDR provider risks slipping.
- As well as tracking the MDR provider's performance against its SLA commitments, the customer organization should also track how well it completes its own actions – such as time taken to actually view the MDR provider's completed investigation

Organizations must define clearly how the new security policies and requirements agreed with the MDR provider are going to be implemented internally according to specific milestones.

Figure 6: State of the Art MDR Overview



Source: HardenStance

Provision also needs to be made to support an MDR provider's threat hunting activity.

- notes. Metrics like these can serve as key inputs into an organization's ongoing investment decisions in technology, services and staffing levels.
- As well as working with an MDR provider to develop and maintain IR playbooks, provision also needs to be made to support an MDR provider's threat hunting activity. By their nature, most threat hunting projects can't be playbook-driven. An organization must ensure that its MDR provider has adequate internal touch points who are authorised to give threat hunters spontaneous access to the internal company resources they need to carry out their work.

Summary

It's getting harder for all but the most well-resourced organizations to manage cyber risk effectively. State of the art SOC capabilities offering visibility and threat detection across endpoint, network, cloud and SaaS assets, depicted again in **Figure 6**, are critical but very challenging to build and operate. For many organizations, partnering an MDR is the best, or only, way to substantially reduce the time they take to detect and respond to cyber threats.

SLAs are critical but buyers shouldn't focus on them exclusively as a basis for selecting an MDR partner. Other factors need to be weighed to assess the MDR provider's long term commitment to the market and how softer, less easily measured, added value can enhance the organization's security posture. Organizations must also recognize that for an MDR partnership to succeed they must have a strong commitment to serving as an enabling partner for their MDR provider - and institutionalizing that commitment throughout their internal processes ■

About the Sponsors

The sponsors of this White Paper are AT&T, ElevenPaths, Fortinet and Secureworks.

About AT&T

AT&T Cybersecurity helps to reduce the complexity and cost of fighting cybercrime. Together, the power of the AT&T network, our SaaS-based solutions with advanced technologies including virtualization and actionable threat intelligence from AT&T Alien Labs™ and the Open Threat Exchange™, and our relationship with more than 40 best-of-breed vendors, all accelerate your response to cybersecurity threats.

Our experienced consultants and SOC analysts help manage your network transformation to reduce cybersecurity risk and overcome the skills gap. Our mission is to be your trusted advisor on your journey to cybersecurity resiliency, making it safer for your business to innovate. Learn more at <https://cybersecurity.att.com/>

About ElevenPaths

At ElevenPaths, Telefónica Cyber Security Company, we believe that a safer digital world is possible. We support our customers with their digital transformation, creating disruptive innovation in cybersecurity in order to provide the necessary privacy and trust in our daily digital lives. We combine the freshness and energy of a start-up with the knowledge, power and strength of a global Telco to provide innovative solutions spanning across prevention, detection and response to daily threats in our digital world. We also work to ensure a safer digital environment through strategic alliances that allow us to enhance our customers' security, as well as through collaborations with leading bodies and entities such as the European Commission, Cyber Threat Alliance, EUROPOL, INCIBE, and the OAS. Learn more at www.elevenpaths.com

About Fortinet

Fortinet (NASDAQ: FTNT) secures the largest enterprise, service provider, and government organizations around the world. Fortinet empowers its customers with intelligent, seamless protection across the expanding attack surface and the power to take on ever-increasing performance requirements of the borderless network - today and into the future. Only the Fortinet Security Fabric architecture can deliver security features without compromise to address the most critical security challenges, whether in networked, application, cloud or mobile environments. Fortinet ranks #1 in the most security appliances shipped world-wide and more than 375,000 customers trust Fortinet to protect their businesses. Learn more at www.fortinet.com

About Secureworks

Secureworks® (NASDAQ: SCWX) is a technology-driven cybersecurity leader that protects organizations in the digitally connected world. Built on proprietary technologies and world-class threat intelligence, the company's applications and solutions help prevent, detect and respond to cyber threats. Red Cloak™ software brings advanced threat analytics to thousands of customers, and the Secureworks Counter Threat Platform™ processes over 300 billion threat events per day. More than 4,000 customers across over 50 countries are protected by Secureworks and are Collectively Smarter. Exponentially Safer.™ Learn more at www.secureworks.com

About HardenStance

HardenStance provides trusted research, analysis and insight in IT and telecom security. HardenStance is a well-known voice in telecom and enterprise security, a leader in custom cyber security research, and a leading publisher of cyber security reports and White Papers. HardenStance is also a strong advocate of industry collaboration in cyber security. HardenStance openly supports the work of key industry associations, organizations and SDOs including NetSecOPEN, AMTSO, OASIS, The Cyber Threat Alliance, The GSM Association and ETSI. To learn more visit www.hardenstance.com