THE SECRETS OF EVALUATING SECURITY PRODUCTS

By
Les Correia
of Estee Lauder

and
Migo Kedem
of SentinelOne

Sponsored by SentinelOne

Table of Contents

Foreword by Phat Hobbit 3
Introduction 6
- Summary 6
- Why we wrote this book 7
- About the authors 7

1. The Modern Challenges of Securing the Enterprise 8
- How cybersecurity evolved 8
- How next-generation security solutions can cope with that 11

2. Plan Your Play 16
- Map your business needs 16
- Be Clear About Your Resources 19
- Understand Internal and External Political Processes 20

3. Success Criteria Validation 23
- Business Management Drivers 23
- Security Drivers 26
- Operational Drivers 27

4. The Evaluation Process 30
- Where To Start 30
- Pitfalls 31
- Testing 32
- Testing Approach and Methodology 33

5. Deployment Preparation and Planning 38
6. Head-to-Head: A Vendor and a CISO on the Evaluation Process 40
7. Conclusion 48
Appendix A – Table - Capability/Feature, Priority, Ranking 49
Appendix B - Malware Testing 57
- Bring Your Own Malware (BYOM) Guidelines 58

Foreword

By
Phat Hobbit

I've been watching the anti-malware industry for a long time as a user, service provider and now cyber security industry analyst. For many years it was the only line of cyber defence and most of the folks I know from "back in the day" have a love hate relationship with their anti-virus provider. Les Correia and Migo Kedem are both heavy weights in the anti-malware industry and their collaborative efforts have produced the book you are reading now. To be honest I'm not even sure folks read introductions anymore as the sentiment is "why bother when it's really the content I am after?"

Fair point, but if you are reading this now, I think it's time to explain just how important this book is from a business perspective. For me I believe we have evolved from a very two-dimensional relationship to cyber space to a three-dimensional cyber environment. Most small & medium business were architected along the lines of "inside the firewall" = things I need to care about and "outside the firewall" = Things I don't need to care about - a very 2-D space. Along came SaaS and cloud hosting and we find ourselves plunged into the deep end of the pool. The third-dimension is really about "someone else's computer that I need to care about". Those in enterprise architecture

have been living in 3-D space for some time with data-centre architecture but there is a new twist: "someone else's computer that I need to care about and have limited control over" 3-D.

Across the spectrum of business no one can argue that SaaS and cloud hosting have become a focus of organizational direction pushing IT departments large and small. This evolution has created immense complications for cyber security ranging from beliefs that "cloud is secure by default" to "we don't have any skill or tools suitable for secure cloud adoption". What I think is universally true is that end-points - specifically user workstations and devices - are the targets of cyber criminals as they provide access to data and resources no matter where they are in cyberspace.

In the 3-D cyber world we have an additional complication which was unanticipated. Those user workstations and devices quite frequently escape the protection of "inside the firewall" as an increasingly mobile workforce accesses data and resources from anywhere that is connected - home, public spaces and work. If protecting data and resources in the cloud was a profound challenge imagine global companies trying to secure end-points located across the planet.

This book seeks to educate and assist in understanding the current and future situation in cyber space and provides advice and council on how to prepare your organization for the future. The cyber-criminal problem will not be vanquished anytime soon. In fact, it may be the highest growth "industry" related to the Internet. In a conversational style and business friendly language Les and Migo provide wisdom and strategy to address the business issue of "24/7/365 connection requires 24/7/365 protection". In major data breach investigations, the culprit is frequently a missed patch, unpatched vulnerability or configuration mistake that facilitates a malicious actor's entry into the network. Imagine a solution that gives you confidence in protection while you take time to test and deploy patches to your end-points?

My belief is that most organizations are now understanding that robust cyber security promotes growth, enhances brand reputation, protects profitability and facilitates growth. This book makes the case that robust end-point defence is an important part of the cyber security strategy to support the organizations strategic goals. Data breach is a stressful and costly endeavour for any organization to endure. Perhaps it's time to read some expert advice on the subject and adopt your end-point defensive strategy accordingly.

Enjoy the Read!

Phat Hobbit

London, *August 2019*

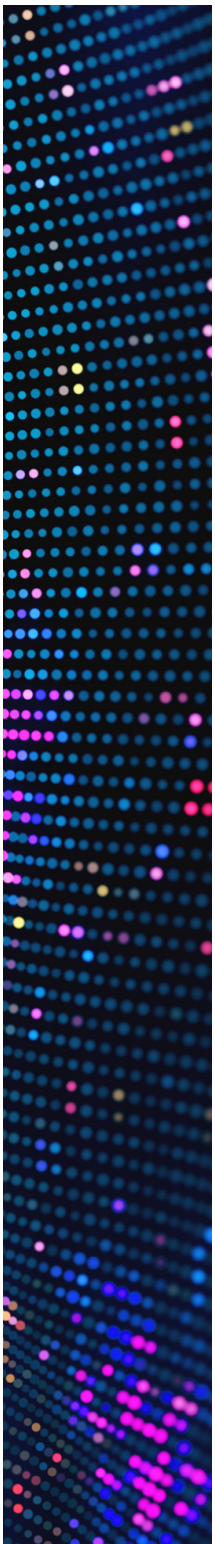

Introduction

Summary

Choosing the right security products to suit your business is a serious problem. The risk of being compromised is real, and there is a lot at stake. The cybersecurity landscape has changed significantly in the past decade. This book puts that into context and explains why modern challenges require next-generation tools.

However, this book is not simply a discussion of trends; it will provide you with a methodology to help you execute the product evaluation process. Starting with mapping your business needs, we explain step-by-step how you should determine those needs and what solutions would work best for you.

Then we break down success criteria validation, looking in detail at drivers from the perspectives of business management, security needs, and operational requirements. The interplay of these drivers must be carefully managed to achieve a good product fit.

A walkthrough of the evaluation process follows, with practical tips on where to start, pitfalls to avoid, and how to create success criteria that suit your business. There are detailed checklists in the appendices that will further support this process.

Finally, we close with a head-to-head discussion between the authors, one who approaches the issue of security tool evaluation from the perspective of a vendor and the other who addresses the process from a CISO's point-of-view. We hope that the resulting discussion adds some enlightening points and provides an enjoyable close to the book.

Why we wrote this book

The challenge of evaluating security products is one that every CISO faces and both of us have faced it many times, from our different perspectives. We wanted to share some of our experiences so that security professionals going through the process have a framework to apply that will simplify the task. The proposed framework may be used to evaluate security products even though we use the evaluation of endpoint security product solutions to expound the context of this basis.

About the authors

Les Correia is a veteran cyber-security thought leader who has successfully chosen and implemented security solutions in leading Fortune 500 companies. Migo Kedem has spent his career building security products in companies like Checkpoint, Palo Alto Networks, and SentinelOne.

Section 1

The Modern Challenges of Securing the Enterprise

How cybersecurity evolved

Cybersecurity technology has become increasingly sophisticated over the last decade. Tools for securing the enterprise are faster and stronger and, as processing speeds have increased, it has become possible to crunch far more data and even apply machine learning to get smart about threat identification.

Life would be simple if just one side were getting stronger in this battle. Unfortunately, that isn't the case, as threats are evolving too. First, because today's attack surface is larger than ever. More devices are inter-connected and sharing data – not just PCs but printers, air-conditioning systems, speakers, lights and even vending machines. Many of these devices are often developed with security as an afterthought. The Internet of Things (IoT) means that this problem will increase. Add to that the rise of Bring-Your-Own-Device (BYOD) and security experts often find themselves trying to secure a flood of employee devices in addition to those of the enterprise.

Second, cybercriminals have access to the same advanced technologies as security experts, and that increases their chance of exploiting vulnerabilities in this broad range of targets. It's relatively trivial for them to try more and more attacks in the hope that just one succeeds. On the other side of the fence, just one failure by a security team can mean disaster. Ethical hackers, researchers, and developers work hard to identify weak points and secure them, but attackers share tools, techniques, and information. In many ways, they collaborate more effectively than security specialists!

Finally, on top of all this are regulatory concerns around privacy and data security. The arrival of the European Union's General Data Protection Regulation (GDPR) in 2018 added complexity to securing data and managing breaches, and the California Consumer Privacy Act (CCPA) is set to bring similar concerns to the U.S.

While all this has been going on, security has broken free of its old home in the IT department. A majority of firms now have dedicated cybersecurity experts, and almost every company now considers security to be "everyone's responsibility," rather than just confined to IT. In practice that means staff should be careful about what emails they open, the files they download and should refrain from installing their own apps and running 'shadow IT.' However successful that is, dedicated security professionals are needed more than ever to assess threats and manage the tools used to tackle them.

Managing risk requires an adaptive and agile security culture, one that binds process, technology, and people together in a way that is effective and that allows the organization to act smarter. Having the right security products in place is essential, of course, but when it comes to adding to your arsenal, how do you know that what you are buying will be effective and worthwhile?

Achieving a mature security culture means embedding the risk and security teams in the process of evaluating products, vendors, and services. A typical procurement process might not be sufficient. Comparing specifications takes time and expertise, and different vendors sometimes use the same

terms to mean different things. Is the latest innovation from your usual vendor actually new or is it a rebranded version of existing technology? Is the new bit of kit from a vendor you haven't used before actually capable of doing all that it promises?

Factoring into this is cost. Are you buying a security solution outright or does it require a subscription? What's practical and affordable depends on the available budget. EMA's Security Megatrends 2019[1] report found that, although IT budgets have been increasing across sectors, there are still some industries that lag, especially manufacturing, healthcare, pharma and medical.

These verticals have been lagging behind other sectors for years, which makes them a target for attackers. Worse, personal health records are especially sought-after and trade for a high price because they can be used for a broad range of crimes, from acquiring new credit cards, making fraudulent purchases or full identity theft. Manufacturing is a similarly tempting target for industrial espionage. Without sufficient budget to tackle every threat, it's vital to thoroughly evaluate every new product to get the most out of your money.

Typically, thorough planning and assessment will lead to a better evaluation of expected risks in decisions related to adopting new technology. We must understand the use case and our users, validate our resources (both in terms of personnel and infrastructure) and capacity, and expect the unexpected.

Furthermore, it is imperative that we shift our cybersecurity strategy from outright prevention, which is unrealistic given the modern threat landscape, to implementing techniques that quickly detect breaches and limit the damage once a violation is confirmed. Resilience and recovery will become differentiators. Intuitive toolsets can go a long way to speeding up that process.

---

1 https://www.sentinelone.com/blog/ema-security-megatrends-2019/

How next-generation security solutions can cope with that

Security technologies are designed to reduce the likelihood, length or scope of a breach. The rapidly-changing threat landscape has made traditional anti-virus (AV) insufficient for securing a modern enterprise. Legacy AV focuses on reducing the likelihood of a breach but, as any security professional will know, it isn't even particularly good at doing that. In 2014, an executive at a leading AV firm acknowledged that its software is about 51 percent effective[2].

Legacy AV was based on detecting malware through signatures – typically a hash of the file – and later through finding tell-tale strings contained in the binary (or executable) using search methodologies such as YARA rules.

In response, malware authors began to sidestep signature-based detection by padding their files with extra bytes or encrypting strings in ways that were harder to read using binary scanning. Meanwhile, attackers had evolved their approach beyond just writing malicious files to a target device. Instead, they were using fileless attacks to exploit built-in applications and processes, compromising networks by phishing for user credentials or simply stealing computing resource for cryptomining. Furthermore, they used admin-type scripting tools, such as PowerShell, and legitimate administration applications, such as PsExec or TeamViewer, to evade detection while taking advantage of the elevated privileges that come with utilities. Legacy AV couldn't deal with this new wave of tactics.

It's no surprise, therefore, that more and more firms moved away from legacy AV in favor of endpoint protection platforms (EPP) that integrated AV, firewalls, etc. However, this approach was still fundamentally based on known malware and techniques (that includes signature-based), so it did not solve the inher-

2 https://www.wsj.com/articles/symantec-develops-new-attack-on-cyberhacking-1399249948

ent problem with legacy AV. The threats from emerging nation-state actors, cyberwarfare and the trading of hacking technologies on the darknet made enterprises realize that they needed visibility.

The next move, therefore, was Endpoint Detection and Response (EDR), which was a term coined by Anton Chuvakin, of the Gartner Blog Network in 2013, as a means of classifying a new group of tools and capabilities that focused on the detection of suspicious activities on endpoints. Instead of seeking to identify specific malware, these tools were different because they looked instead for anomalous activity. Rather than identifying and quarantining a suspect file, they would alert the security team that there was something requiring further investigation.

No solution is perfect, and EDR is not without drawbacks. Some might claim that attempting to provide the enterprise with visibility of what is occurring on the network is an easier nut to crack than quarantining threats because it shifts the emphasis onto the human agent who has to respond to the alerts generated by the system. Increased visibility means an increased amount of data, and consequently an increased amount of analysis.

This means that most of today's EDR solutions aren't scalable. They require too many scarce resources – time, money, bandwidth, and a skilled workforce. Also, because today's EDR requires cloud connectivity, it will always be slower than a solution that is on the device. A successful attack can compromise a machine, exfiltrate or encrypt data and remove traces of itself in fractions of a second. Waiting for a response from the cloud or a human analyst is simply not feasible when dealing with those types of threat.

	Traditional AV	Next Genaration Toolsets
Goals	• Prevent Attacks	• Prevent • Detect • Respond • Analyze
Methods Used	• Virus Signatures/Known Malware • File Analysis • URL blocking • OS Behavior Analysis • Web/File/Source/Reputation • HIPS Malware – Antivirus using signature based detection IoT & Endpoints – Manual device-level security updates though management station on premise or the cloud	• File and Fileless Execution Analysis • Command and Control Detection • Ransomware Protection • Behavior Analysis • URL Blocking • Artificial Intelligence/Machine Learning • Deep File Analysis • Auto Remediation • Root Cause Analysis and First Level SOC Malware – Pattern recognition and predictive analytics, learning to thwart new attacks IoT & Endpoints – Device and network behavior analytocs, anomaly detection
Malware Type	• Previously Known Virus/Malware • Static Antivirus Threats	• Previously Know Virus/Malware • New Morphing and Fileless Malware, Ransomware

Solutions	• Traditional Antivirus Suites	• Next Generation Endpoint Security
		  • Endpoint Protection and
		  • Endpoint Detect and Response
		
		• Integration with CASB, Gateways, Firewalls, Security Solutions

Figure 1- Traditional Endpoint AV versus NG Endpoint Toolsets

Next-generation solutions evolved as a result of the above-mentioned limitations of EPP and EDR. Next-generation products are attractive and can be lightweight because they combine data science, threat intelligence, artificial intelligence (AI), machine learning (ML), and cloud analytics to sift through and analyze associations between patterns of behavior to detect the tactics, techniques, and procedures (TTPs) used by attackers. Patterns of malicious activity and behavior can then be identified and exposed through analysis and correlation. Additionally, they typically use scalable consoles and storage, which alleviates maintenance and other costs.

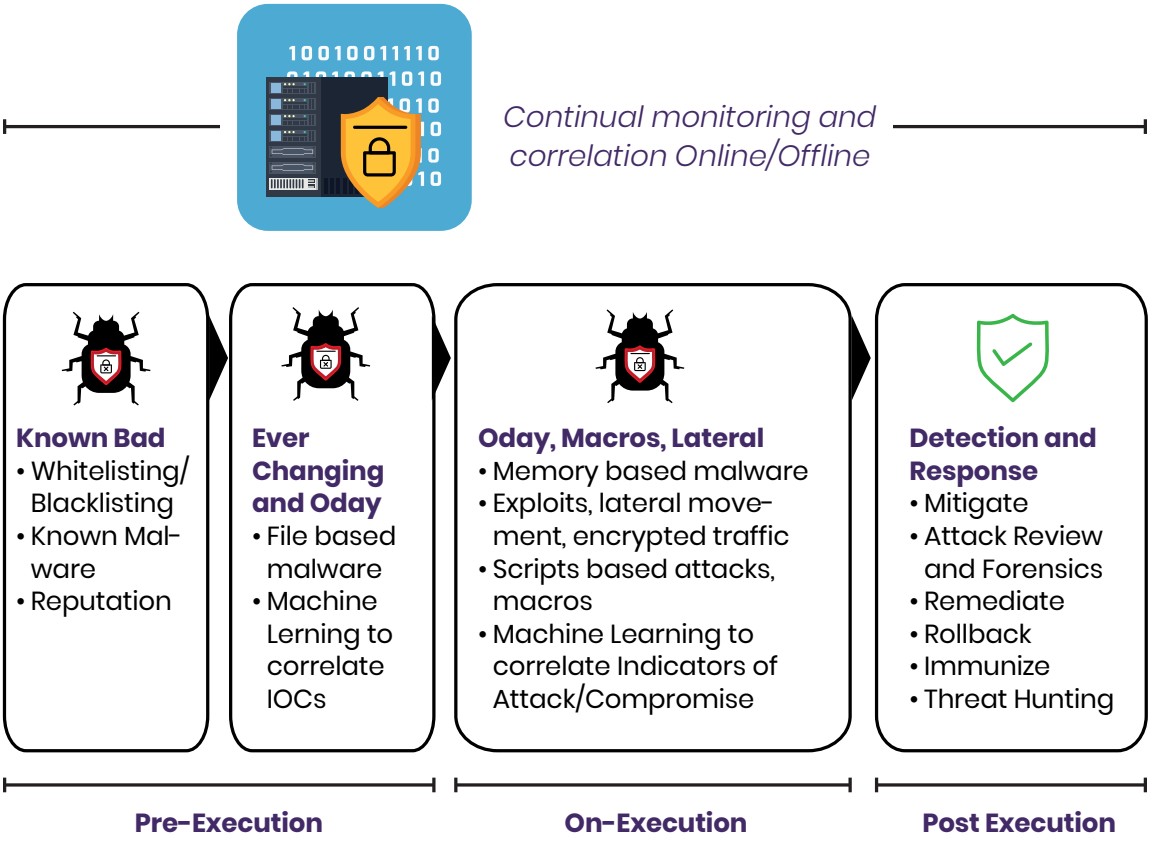

Figure 2- Generic Next Generation EPP Flow

Next-generation EPP solutions use layered protection by correlating and scoring particular behaviors and indicators of compromise. Different vendor solutions demonstrate strengths in certain areas, for example, a specific phase of an attack. Endpoint agents apply critical thresholds and allow the activity to proceed through a 'gate' if the threshold is not met.

Almost all endpoint agents use a reputation phase-based test, as in traditional AV. If the malware does not meet the threshold and therefore passes through the gate, then it will next be tested statistically, pre-execution, using AI to characterize it. The malware indicators are then added to the reputation score. If the reputation score reaches the threshold, then it is flagged; otherwise, it passes to execution.

Next is the execution phase. Most products defer, that is, some allow for full execution and block specific patterns, others execute in a sandbox, perhaps in the cloud, and some do not get to this phase at all, possibly erring on the side of caution and raising a false positive at the previous stage.

EPP that stop any one of the aspects of an attack would either completely prevent the attack or significantly reduce the damage. An EPP product that can stop all the aspects is more effective than one that depends upon a single protection method to do all the work. Additionally, a high number of false positives is a good indicator of an overly sensitive product and one that often does not have the depth of functionality and capability to clearly distinguish good from bad files/behaviors.

Some next-generation endpoint products add a post-detection phase that includes Endpoint Detection and Response (EDR) that allows for incident response. The latest approach is next-generation EDR that uses AI to provide an automated response that takes the burden off the security operations/analyst team. It allows security teams to quickly understand the story and the root cause of a threat.

These next-generation EPP and EDR tools can construct a storyline – a chain of events – from which they can learn the appropriate lesson and store that, and variations of it, for the future. Like any AI or ML process, the effectiveness of the process comes not so much from the algorithm used to analyze data but the quality of the data itself. How the data is collected, analyzed and classified is as important as the algorithm itself. This is an important consideration when evaluating efficacy. Another critical factor is whether EPP and EDR are combined or separate products.

Section 2

Plan Your Play

Map your business needs

Before you even begin to evaluate any security toolset, you need to start by deciding what you are looking for – your requirements. That means, first, mapping-out your business needs, then assessing your available resources and finally, take a clear view of the internal and external political processes that might affect your decision-making and roll-out.

In general, there is a lack of standardization in testing security products to validate their efficacy against the broad spectrum of rapidly evolving threats. Without planning, many organizations must rely on third-party tests, vendor-prejudice, or simply marketing and sales material.

When it comes to evaluating security products, a meaningful evaluation is impossible without clear success criteria. Success means different things to different organizations. For one, it might be a priority to support legacy devices because they have a long tail of unpatched devices that will remain in use for some time. Another organization might need to ensure macOS security, while a third might consider it most important to have a solution that its current IT workforce can manage easily. Unless you understand your particular needs, then you won't be able to determine whether a specific product is a good fit in the first place. And, later, you will be unable to establish whether it is doing what you need it to do.

It's easy to fall into the trap of assuming that a breach won't happen to you. However, Security Megatrends 2019 found that 73 percent of respondents had been affected by some form of endpoint attack and only 58 percent of orga-

nizations were highly confident that they could detect an important security incident before it caused a significant impact.

A good first step for assessing your success criteria is to go back to basics. What are you trying to protect? For most businesses, there is some information that is critical and must be protected at all costs – the 'Crown Jewels.' This could be patent (or trade secret) information or banking information for high net-worth individuals. Whatever it is, it probably requires extra security.

Then there is the rest of the IT estate. Do you know where the boundaries are? Which devices are connected to the network but perhaps don't get used very often or might not be patched as they should? Is there a 'shadow IT' problem in your organization, where staff members have sourced their own devices or applications without informing IT? That might also pose a threat. And, talking of users, do you know all of the devices they use to connect to the internet? Most organizations audit these things regularly, so you might not need to do the above if you have a recent audit to look at instead.

The obvious next question is: who are you trying to protect yourself from? Because different organizations have distinctive types of data to protect, they may attract contrasting threats. For many companies, the biggest risk is an insider threat. This is often unintentional or negligent – an employee accidentally makes some information public that shouldn't be, or they respond to a phishing email. This can be handled by policies and training. Occasionally, it's the result of a malicious employee, someone who feels mistreated by the company or holds a grudge against a colleague. This is could be handled by micro-segmentation, deceptive technologies or similar solutions.

Anything that isn't an internal threat is, by definition, an external one, though that takes different forms. Nation-state actors are seeking to disrupt business as usual or steal secrets, cybercriminals whose primary goal is profit, and threats from partners or third parties who have their own insider risk.

What these attackers want matters because it affects how you might detect them. If you are concerned about hackers who want to steal data, then you would need to look for signs of data exfiltration, among other things. Ransomware, a relatively new threat, is about merely locking machines and demanding payment to unlock them, so determining that is about finding people with unauthorized access.

As mentioned in the previous chapter, we have seen a rise in so-called 'fileless' attacks recently, and these are specifically designed to avoid triggering the systems that are in place to detect intrusion to the network. SentinelOne's H1 2018 Enterprise Risk Index Report found that fileless-based attacks rose by 94 percent, compared to the previous year[3].

Fileless-based attacks use system files to run malicious code, for example by launching attacks against a running system process such as iexplore.exe or javaw.exe, and thereby avoiding leaving a footprint on the storage system that an endpoint agent or file integrity monitoring tool might catch. One of the reasons that 'fileless' malware is so useful to attackers is that security products cannot just block the system files or software that are used in these attacks. If a security admin blocked PowerShell, for example, then IT maintenance would suffer. The same is true of blocking Office documents or macros, which attackers can also use to trigger malware.

Also more common are cryptojacking attacks. Mining cryptocurrencies have become increasingly expensive as the amount of computing power needed is costly both in terms of equipment and electricity. Hijacking other people's machines and using them to mine cryptocurrencies, with the proceeds sent to your encrypted wallet - this new trend in malware has subsided slightly as cryptocurrency values fell in the recent past.

---

3 https://www.sentinelone.com/press/sentinelone-unveils-h1-2018-enterprise-risk-index-report/

Assessing your success criteria means auditing your current tools and looking for gaps. Where are your current tools weakest and where do they not cover you at all? How have the threats changed since you bought those tools? What emerging threats do you expect to have to deal with and how well do your existing tools protect you from those?

Be Clear About Your Resources

We mentioned in the previous chapter that EDR solutions do a great job of finding potential threats, but that also means that they generate a large number of alerts. Somebody has to investigate those alerts and determine whether they require action. Few companies have lots of analysts available to pick up the alerts. When evaluating a new security tool, it's important to consider whether it requires more staff than you have available.

Even if you do have sufficient staff, the work that they are doing now might not be the most effective use of their time and skills. It may be that bringing in a next-generation security solution can free-up some of your current team to do more useful tasks than merely responding to alerts. In most security teams, it is easier to automate repetitive tasks than it is to increase the headcount.

The headcount problem is exacerbated by the ongoing shortage of staff with cybersecurity skills. The number of organizations reporting a cybersecurity skills shortage has risen every year from 42 percent in 2015 to 53 percent last year. Research suggests that there will be two million unfilled cybersecurity positions by the end of 2019 and that will almost double – to 3.5 million – by 2021[4].

How many skilled staff you have will affect your determination of whether a particular security solution fits your needs. A further consideration is whether

4 https://cybersecurityventures.com/jobs/

you expect staff numbers to stay roughly the same, to increase as the company grows or to shrink with potential cuts. If it's likely to increase, how confident are you that you could fill those positions? If you expect your team's workload to increase but numbers to stay the same, then this should also factor into your evaluation.

This is an area where technology can help, and it could well be the reason that many organizations seek new or upgraded security products. Tools that are smarter or easier to use require less specialist knowledge on the part of the humans working with them. The machine automates many of the complex parts of the task, making security less labor-intensive and processing far greater amounts of information than humans can.

Another factor to consider is the level of expertise that you have across your team. Do they have the skills necessary to manage the new system? Even better, is the new system compatible with their current skill level, meaning that they can adapt it with minimal training and adjustment?

While interviewers like to ask candidates 'where do you see yourself in five years' time?', it's good to ask the same about your current cybersecurity defenses and strategies. How does the new tool fit not only with your existing tools but also with tools that you are likely to adopt in future? And, perhaps even more importantly, how does this plan fit with your expectation of future staffing levels? Lastly, how do these tools integrate and share intelligence to produce consistent feedback?

Understand Internal and External Political Processes

Every company has a security culture and understanding yours is crucial when evaluating security products. Politics plays into the decision making in several ways. First, it affects staffing. A company that values the security

function and understands the risk of a breach is likely to fund it accordingly. A company that sees it as a cost to be trimmed as much as possible will be less willing to spend. Most companies will be somewhere in between, with security valued but competing for funds with other parts of the business. This affects not only the kind of product that you can afford but also the medium and long-term view of how that product might be supported.

Second, a corporate culture that sees security as everybody's responsibility and actually takes steps to embed that in behavior (rather than merely pay it lip service) is likely to put less of a burden on its security team. There should be fewer avoidable mistakes and less need for the team to be continually educating people. That means more time devoted to handling significant alerts and focusing on deeper threats.

Third, in some organizations, the CISO reports directly to the CEO and the C-Suite executives are very focused on how security is handled across the business. Other companies position security as a sub-department of IT, wherein influencing the overall culture of the organization is more difficult. A stand-alone security team can argue for its own funding, without having to argue for a portion of the overall IT budget, and is better able to make a case for itself directly to the C-Suite. Furthermore, this can enhance visibility at the board level.

External political processes also need some consideration. This includes capital-P Politics – local, national and international government – because your organization might be considered critical infrastructure and therefore have specific security responsibilities, or it might be a regulated industry and therefore required to attain certain standards. Having a weather eye on possible changes in legislation and new requirements will help you to determine whether the tool you're evaluating will help you to meet obligations that are not yet official.

The need to consider external politics also refers to politics with a small P – relationships with suppliers and partners and how these might have an impact

on security considerations. Perhaps you are connected to a third-party with security habits that are less than ideal. For political reasons, it might be beyond your remit as CISO to end this relationship, but it's certainly necessary for you to do what you can to secure it. Again, being aware of political sensitivities around such relationships will factor into your thinking when it comes to adding new security tools.

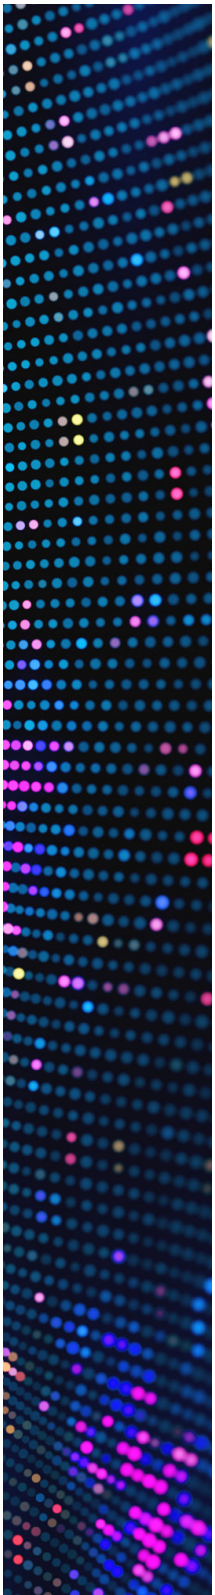

Section 3

Success Criteria Validation

So far, we've considered the changing threat landscape and how security tools have changed to keep up with that, and we've looked at some of the things you should consider when determining what kind of product you plan to look for. The next move is to understand what success looks like. We can divide this into three areas: Business Management, Security and Operational Drivers. Let's consider each set of criteria in turn.

What follows is by no means an exhaustive list. There might be some that apply to your organization that we have not covered here. If so, you should include those when making your own list and then assign each one a High, Medium or Low priority. This is important because a product that ticks a lot of boxes is great but if those boxes are mostly your low priorities then the product is leaving out some important needs. No tool is perfect, so there are likely to be some areas where you might need to compromise. This section will help you to determine where those might be.

Business Management Drivers

A good place to start is by considering the business management drivers for your purchasing decisions. The new tool will need to fit your governance rules, processes and resource capacity but it will also be intended to help fix some problems or make some things run more smoothly. Assessing all of this and attaching some priorities will help you in your evaluation process.

1. Does the tool comply with your organization's regulatory requirements? This is likely to be a high priority for every organization, and it might require scrutiny of the small print. If the tool you are evaluating sends some data to the cloud, for example, then you might need to consider where the vendor's servers are. In some regulated sectors, there are strict controls over whether – and what kinds of – data is allowed to cross borders.

2. The product you are considering might meet your needs today but does the product's road map fit with your organizational needs and how you expect them to develop in the future? The road map might suggest that the vendor intends to focus on areas of the product that are not relevant to you, for example, or the road map itself might not be as clear as you would like. Things change quickly in the cybersecurity world, but you need to consider how you expect the next five years to unfold and determine whether this product fits your vision.

3. Do you anticipate that the product you are evaluating will produce fewer alerts? This might be because it is better at eliminating false positives or because it is more automated and will handle specific routine alerts by itself. Either way, this will deliver an improved experience for end-users, which is likely to enhance job satisfaction and might even have an effect on employee retention.

4. Will the new product allow you to simplify your product suite by consolidating solutions? If so, is that a significant driver? Some organizations prioritize a simple suite of products, and others are happy to use as many products as it takes to feel secure. Most are somewhere in between.

5. The vendor's deployment and licensing policies will affect new features on the road map. Don't assume that you will automatically and instantly get access to every new feature as it is released. Check whether you will have to wait for new features or if access to new features depends on the terms of your license. How much the answer matters will depend on the weighting, you give to new features generally and links to point 2, above.

6. Consider the vendor's background. If it's a new vendor, then check references, reviews and whether analysts have assessed the product. If it's a longstanding vendor, then consider the company's track record in terms of product success, company stability, and customer service. A well-funded new startup that everyone is talking about might have a great brand-new solution, or it might not. The company might run out of funds or could be acquired and change direction.

7. What is the vendor's support provision like and what is its approach to SLAs? How you weight this consideration will depend on factors like the size of your team. If you have a small team, then you might need to rely on support much more. Similarly, the business you are in will have an impact on the kind of SLAs you expect.

8. What is the product documentation like and is it regularly updated? Again, how you weight this will depend on whether you expect your team to need to refer to the documentation frequently or only during setup. It will also depend on whether your team is already familiar with similar products.

9. How much weight do you place on references? As mentioned in point 7, above, references might be more important for a new company than one with a long history, but you will still want to get references from companies using the solution.

10. Finally, what financial drivers are behind your purchase? Is there pressure to consolidate tools or to free-up staff resource, for example? Or are economic concerns secondary to getting the right product? A proper return-on-investment calculation will go a long way to set expectations rather than assumptions.

Security Drivers

Next, you need to consider the security needs that are driving your buying decision and assign them a priority. As discussed earlier, not every product will meet every demand and all products will meet some needs better than others.

1. First, then, how important is protection and detection? Are you looking for a tool that will reduce your exposure to known and unknown attacks, whether from malware, exploits or blended threats?

2. To what extent are you looking for something that is future-proofed? You might expect extensible attack intelligence and analytics, with automated updates that protect against new attacks. On the other hand, perhaps you are looking for a tool that will do one job, and you do not expect that job to change very much.

3. How significant are the risks of intellectual property theft, reputational damage or privacy invasion? Some tools will protect against these risks better than others.

4. Are you expecting efficient EDR? This ties in to point 4 in the previous "Business Drivers" section. An efficient, next-generation EDR solution will produce less security 'din' than legacy solutions. That means fewer alerts and more automation, which is especially crucial for a team that is small and/or overstretched.

5. Is it important that this tool provides online or offline protection or both? Some tools are stronger in one area than the other. Also, you might have offline protection covered already and only need online, or vice versa.

6. How easy is it to collect forensic evidence from the new tool and how important is that feature? Depending on what you plan to use the tool for, in some cases, this could be vital, while in others it will be irrelevant.

7. Spend some time thinking about interfaces and interoperability. Examine how easily you can collect audit logs from the new tool and how well that integrates with your current log aggregator. Consider the availability of APIs, whether they are sufficient for your needs and how you will use them.

8. How does the new tool help you with user and device identification for determining Patient Zero in the event of a breach?

9. Finally, does the new security solution provide level-1 SOC capability?

Operational Drivers

The last set of drivers to consider are the operational ones. This is the most practical part of the success criteria validation and the one where you should pay attention to how the new security solution will work in everyday use.

1. First, consider the endpoint platform's coverage. Assess the broad operating system and platform support, ease of rollout and/or update – including whether and how you will have to manage reboots – and then the new tool's compatibility with your current environment. You don't want to discover at the last minute that this new tool does not work well with another tool that you use and rely on a lot.

2. Next, look at the enterprise management tools and processes. Is the new tool easy to administrate, with features to assist with alerts, remediation, remote access and so on? What is the response and remediation capability of the new tool? Does it have centralized administration but still allow for regional hubs? Are the dashboards intuitive? Does it have a cloud-

based console, meaning there is no equipment to maintain? Consider each of these aspects and how important they are for your needs.

3. Visibility and contextual awareness are an essential part of any next-generation security solution but consider the extent to which the tool that you are evaluating provides these. Do you get the visibility that you need? How would increased contextual awareness help you?

4. Will the new tool deliver a reduction in hardware and software maintenance? If not, do you have the resources to maintain the tool? If you can expect a decrease in support, then you have extra capacity that you can plan to devote to other tasks.

5. What are the endpoint management options in the new security tool? As mentioned in the first chapter, this is an increasingly vital area to pay attention to how this might affect your needs in the future.

6. Related to the previous point, consider what kind of local endpoint intelligence you need and whether the service you are evaluating meets that need.

7. One of your drivers might well be to reduce alerts from your security tools. However, when those alert communications come through, are they delivered securely? Again, you need to assign a High, Medium or Low priority to that.

8. Spend some time thinking about performance. Lightweight agents will be less of a drain on your resources elsewhere. Determine how much memory usage and bandwidth the new tool requires and consider ease of deployment.

9. In the Business Management Drivers section, above, we noted that regulatory concerns will affect your evaluation of any new security tool. Related to that, you will need to ensure that the new tool provides the kind of reporting capabilities you need for compliance purposes.

10. If the new service consolidates tools or features that you have been using in other systems, then remember to factor in potential training efficiencies. Training your team on fewer systems will save time and money, as well as lightening their cognitive load.

11. Finally, how responsive is the support team of the vendor that you are considering? This ties-in to the considerations about SLAs and support availability mentioned above. Or do they have a complementary managed service that can be introduced to ease support for sophisticated events and dependencies? Be realistic about your needs and expectations of the vendor.

At the end of that process, you should have a list of criteria, each of which has been assigned a priority and the next step will be to cross-reference this list with what the vendor offers. The points above give plenty of examples of what to consider when doing this. You should now be in a much better position to determine whether you want to proceed with the actual evaluation.

Section 4

The Evaluation Process

Where To Start

By this point, you have established that you have a need, you have drilled down into how the tool that fills that need fits with your existing processes, resources and the internal and external politics affecting your implementation. You have also assembled detailed and prioritized success criteria that you will use to assess the new product.

You might already have a particular product in mind, perhaps because of a recommendation from a colleague or because a current vendor has released something new. Possibly you need to put together a shortlist of products that might meet your needs. Either way, it's worth surveying the market to get a sense of how many products might fit your requirements. Some of these will be eliminated quickly – some guidance provided in see section 9 in Appendix A – Table - Capability/Feature, Priority, Ranking.

The cybersecurity market includes many solutions designed to solve different problems, but some vendors are tempted to promise more than they can deliver to gain an edge. The result is that too many products claim to protect organizations from any threat, any attack vector, and without any impact on users. It goes without saying that you need to look behind these claims. It is imperative that you validate your vendor or any third-party testing report. This is not to say that the analysts are wrong but only your test can replicate your production environment and its idiosyncrasies.

Pitfalls

As legacy AV software proved ineffective in recent years, some companies responded by cutting security budgets. After all, if something doesn't work reliably, then why pay all that money for it? In the era of zero-trust networks, there are still cases where the budget for network security is higher than for endpoint protection. Nobody wants to waste money on an ineffective solution. An effective solution secures your most important assets and lets you sleep soundly at night but doesn't need to be prohibitively expensive. Do your research and be sure to allow sufficient budget for what you need.

We mentioned in the previous section that no cybersecurity solution would protect you from every risk, so you are going to need more than one product to stay safe. Network security, emails, access control, logging, and security information and event management (SIEM) usually live under the same roof. If your security products don't talk to each other and cannot contribute to each other's efficiency, then that is a bad sign. When evaluating a new solution ensure that it can integrate with all parts of your security apparatus.

For example, when your firewall gateways determine that a file is malicious, you want to immunize your endpoints against it because they are not always behind the firewall. Likewise, when the endpoint solution identifies a bad URL, you want to block it across your gateway protection. Similarly, if malicious activity is detected, you want to consolidate all the related information from all of your security products to ease the process of identifying the impact and implementing the necessary steps. For that to happen, all your solutions need APIs that allow them to be automated and integrated. Responsible vendors build the necessary APIs as they are building the product, which should ensure that they function as intended. Beware of vendors that build a product, see it mature and then scramble to implement API calls at the request of customers. This approach frequently leads to bugs.

Testing

Feasibility tests usually involve some malware testing, comprising test criteria and a review by the organization's IT and cyber experts. The cybersecurity experts will test their most advanced scenarios but remember that the solution you choose will be enabled on all your devices, for all your users, all the time. While protection is the most critical concern for endpoint security products, your users must be able to work with it without interruption so that they can do their jobs.

While a test may look impressive when everything is set to trigger alerts by default, will your ordinary users be able to work with these settings every day? Many products can protect you when they are set on their strictest level. Anything new, running from the downloads folder or unsigned can simply be blocked. Such strict policies can hurt productivity.

Furthermore, we cannot assume that more alerts mean more success. A solution that overwhelms you with hundreds of alerts is barely usable, and it's certainly not effective. Ask Target, which suffered a major breach when an alert was buried in a pile, where it remained undetected for more than 120 days. You want your team to know when a real threat has been found, rather than be overwhelmed by sheer noise. Detection can be cheap, but it often comes at the price of prevention, as your IT team struggles to piece together the whole story.

When you test a security product, you need to test against an unknown, as well as known, malware. Any security product can run a query against public repositories such as VirusTotal and return those results as its own, so detection of known malware is a low barrier for a product to pass. Testing unknown malware will give you a sense of how this solution will handle novel threats that it might face in the future. The challenge is finding such malware to test. One way to do it is to take a known piece of malware and modify its hash, which will help you to differentiate between solutions that rely on signatures and those that rely on behavioral AI.

As well as file-based malware, you need to include fileless attacks in your testing. As we have already mentioned, these are becoming increasingly prevalent, and you will need to be confident that you have a tool that can detect them.

Another mistake to avoid is assuming that a candidate solution will recognize all your company software as safe. Let your IT users run a candidate solution on their devices as well as on your own. It is not uncommon for large enterprises to run their own software, sometimes even signed by a valid certificate authority. The caveat is that badly written code can exhibit malware behaviors. Doing that will allow you to see a solution's real worth. Nobody else is as knowledgeable about your network as your own employees.

Devise a comprehensive testing plan before you start and be sure to test on users outside of your IT department and ask for their feedback. Once you have completed a thorough testing phase, you will have a much clearer idea of which of your candidate solutions is the best fit for your needs and for your organization.

Testing Approach and Methodology

Before you begin testing, make sure that you have a good plan in place. Set clear objectives, timelines, and deliverables. Ensure that you are prepared to document every step and record the evidence that you collect. Define your priorities High - 2, Medium - 1 and Low – 1, as outlined above, and plan to score each product for each priority on a six-point scale: 5 – Excellent, 4 – Good, 3 – Average, 2 – Fair, 1 – Poor, 0 – Non-existent. You can modify this scale as needed.

Do your research to draw up a shortlist of product solutions that might fit your needs. Aim to have three-to-five products on your list. Too few will limit your options unnecessarily, and too many will result in an evaluation process that

takes too long. We provide some guidance in section 9 in Appendix **A – Table - Capability/Feature, Priority, Ranking.**

Once you have a clear list of criteria and a shortlist of products to test, then you need to prepare your malware testing (see the section below). All solutions need to be tested on the same malware and the same TTPs. Aim to keep timing near real-time to minimize the risk of the zero-day malware you are using being discovered by aggregation sites.

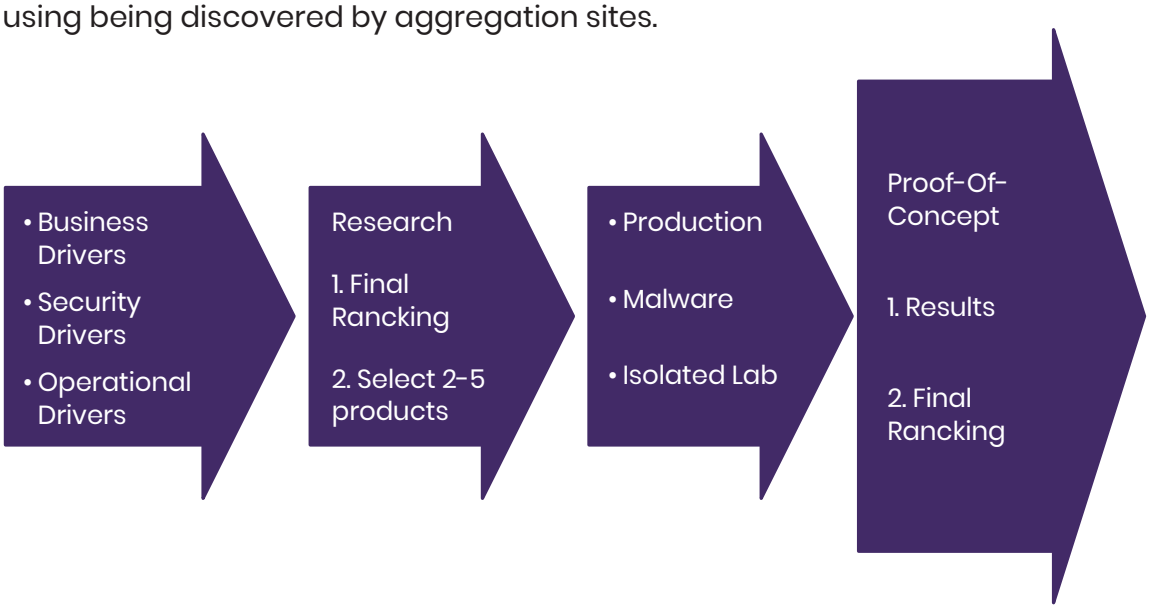

Figure 3 - Proof-Of-Concept Testing approach and methodology

Your testing should compare malware – fileless, ransomware, macros, pre-execution and on-execution, sources for malware, TTPs, and reporting of capture evidence. Test separately in two environments – isolated lab and production – to test product solutions against your organization's business, security and operational drivers.

In the isolated lab environment, the focus is on functional and malware, exploits and blended threats. Ensure documentation and testing for the following threats in the isolated lab in a systematic way:

1. False positives – this to ensure user experience for legitimate software
    - » Use whitelisting for any wrongly detected files to stop alerts. Note

how long it takes to train and if there is an acceptable rate of false positives. You may also be able to find this out from VirusTotal. Is there a way to administer this false positive concern on the console to reduce the impact?

2. Documents and scripts – Weaponized documents, Macros, Malicious scripts
    - Requires careful review on whether the attack was blocked and how it was done. You may have to reach out to the vendor for clarification or read the detail in the documentation. Try executing without scanning or while the victim devices are offline.
3. Executables – Malware, Packed files, Potentially unwanted applications (PUA)
    - Test the victim devices online and offline – individually or en masse.
    - Test packed files and ensure that you pack malicious and benign files.
4. Ransomware variations – Disk and File encryptors
    - Test ransomware on the victim devices online and offline.
5. Exploits – using exploit-based attacks via Metasploit
    - Test exploits by ensuring that the victim device(s) are not being scanned. Also, test when the victim devices are online and offline.
6. Miscellaneous attacks – credential theft, privilege escalation, code caves, etc.
    - Test using the Metasploit Framework.
    - Use internal data available from phishing, previously quarantined, etc.
    - Test for visibility of IoT devices and communications - to confirm threat vectors.
    - Perform a test of victim devices online and offline.
    - Perform test wherein malware downloaded on a USB and tested over several days.
    - Validate lateral movement and notification. Get information on how this is being done.

The lab must be isolated so as not to contaminate production. Furthermore, the use of virtual machines allows for systematic testing and resetting to a

clean snapshot. You can obtain test files from various websites (see Appendix B - **Malware Testing**). The purpose of this testing is to evaluate the features touted by vendors for their solutions.

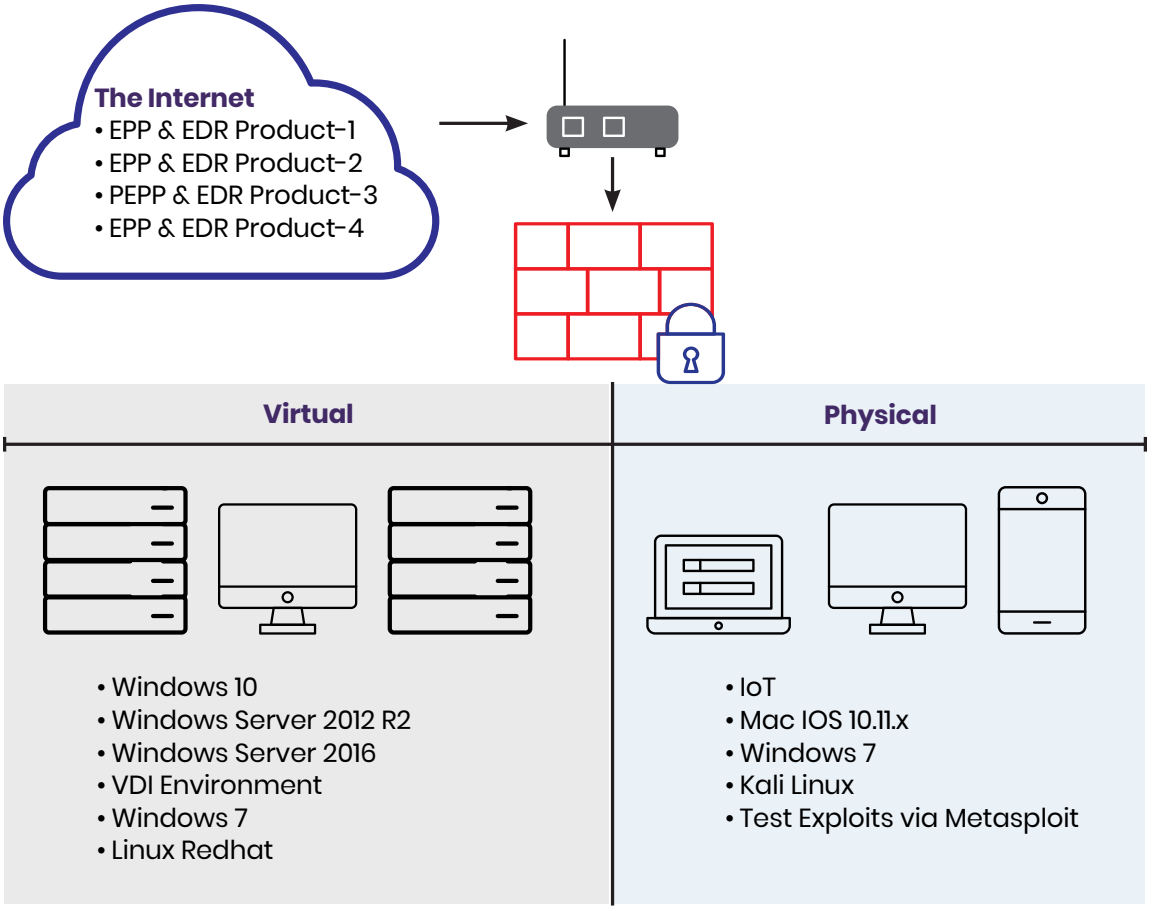

Figure 4 - Isolated Malware Testing Lab

When it comes to testing in production, on the other hand, the purpose is to test compatibility in your organizational setting based on OS images, applications and the tools you already have in use. This includes servers, VDI, cloud containers, and various devices, operating systems and workstations. Follow your organization's processes to get an understanding of how the product will work within your environment.

Throughout your testing compare your observations to the vendor's reference responses. Evaluate from the perspective of your users – endpoint users, application owners, and administrators – and establish possible use cases and

evaluation objectives. Be sure to get feedback from your operational teams. We provide some guidance - see section 9 in **Appendix A – Table - Capability/Feature, Priority, Ranking.**

Once the Proof-Of-Concept is complete, provide a detailed report of your findings. The good news is that you will have the evidence for your analysis if requested.

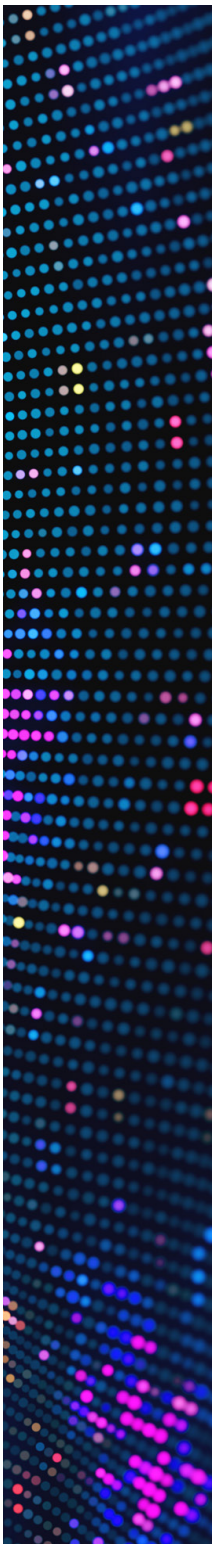

Section 5

Deployment Preparation and Planning

Once the Proof-Of-Concept is complete, plan for deployment. Preparing for EPP/EDR deployment is critical to the successful execution of the project no matter how good your selection. Realize that whenever EPP/EDR touches endpoints, you will discover something new that can potentially negatively impact perception and accomplishment.

Planning best practices:

- Plan the configuration and integration from the EPP/EDR.
  - Would you be utilizing an IAM or Active Directory? How will it be integrated?
  - The Responsible, Accountable, Consulted, Informed, Support personnel? Who requires Read-only access etc.? Confirm the teams.
  - What are the other obligatory tools expected to be integrated? E.g., SIEM, SSO, CASB
  - How do you plan to break down the console? By sites, regions?
  - What policies? And what is the breakdown and how do you expect to enforce?
  - Do you have a set of whitelists/exclusions that are well known? Perhaps from the current EPP solution or third-party applications?
  - Be cognizant about any applications that are I/O intensive, databases, logs, monitoring, and file sharing. They may have to be closely monitored during deployment.

- Validate the processes for alerts, action, handover, and escalations (internally and externally with EPP/EDR vendor) are clear and documented.
- Confirm expectations with the vendor of EPP/EDR.
- Confirm endpoint and any vendor system and installation requirements (software/hardware).
- Validate reports required for any compliance needs.
- Validate documentation pertinent to your environment is available.
- Ensure knowledge transfer/training has been completed.
- Communications with your service desk, business, etc.
    - Any scripts for calls into service desk?
- Build packages for various OS platforms in your environment.
- Confirm strategy for deployment.
    - Perhaps set in monitor mode and then into full protection mode – or similar if possible. This will allow you to see the alerts for a few days/weeks before enforcement. Confirm timeline for this mode change.
    - Also segments of champions for the pilot – e.g, perhaps start with the IT department and then move to others.
    - Maintenance and troubleshooting - Automatic updates? Clarity on disabling, uninstallation.
    - Project-related:
        - Risks
        - Reporting metrics of progress and issues.
        - Issues
        - Change control
        - On-going lessons learned
        - How whitelists, exclusions will be documented and approved?
    - Licenses.
    - Removal/Decommissioning of the incumbent (previous) EPP/EDR.
- Select a pilot environment on a segment of the production environment.
- Confirm that the procedures your team learns (or learned during the proof-of-concept) during the deployment can be leveraged for the entire user base. This will help you customize a set of deployment steps to ensure smooth implementation of your EPP/EDR solution.

- Expand the deployment to several pilot projects on a segment of the production network, all the while training and getting the input of those who will ultimately be managing the EPP/EDR system.

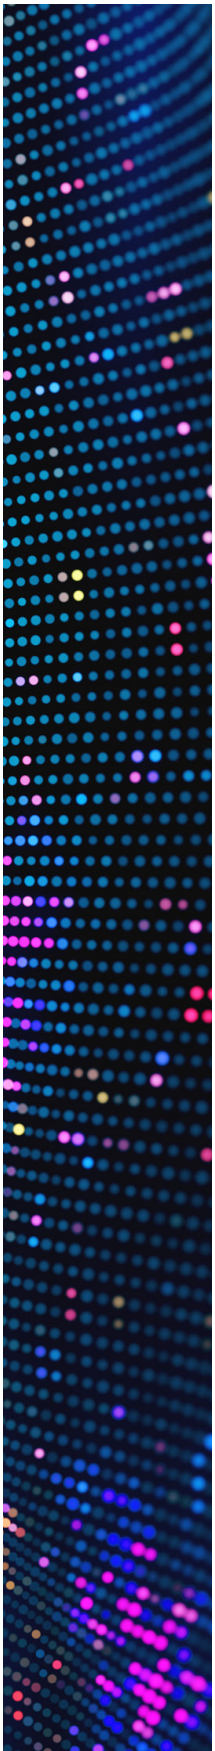

Section 6

Head-to-Head: A Vendor and a CISO on the Evaluation Process

Les Correia

of Estee Lauder

Migo Kedem

of SentinelOne

Les: What's the most important factor for an enterprise in choosing a vendor? In my opinion, one of the key things is to make sure that the vendor's roadmap and the path forward matches the customer's vision and roadmap. It's not so much just the product itself, it's very much if it matches up with what you have but also delivers what you need. Then you can consider other things. How do they integrate with the security stack, and so on?

Migo: There is no one solution that is suitable for every type of organization. It really depends on the problem that the customer wants to solve. It's essential to know your goals for implementing a product. We still occasionally find prospects who are not really sure about what they need. I think it's useful for the vendor to know your goals too so that they can answer those needs. It also helps to ensure a successful relationship.

There is a lot of hype about EDR but most EDR solutions out there require the proper skills and personnel to actually manage the product. When companies don't have that, they end up with a kind of misalignment of expectations, I would say.

Les: Yes, in fact, I was thinking about that 'one size fits all' idea because some companies cannot even be described as a single company. Often, they are a group of companies - a conglomerate - and each of them has a different way of approaching things. You can't approach it saying "I'm buying this product because it will solve the whole thing". One needs to understand the landscape and intricate requirements.

One thing I've learned over the years is that you have to see how responsive the vendor is and whether they adapt. It's also critical to understand how the vendor approaches solving problems. For example, do they hold all the information so that they can expand upon the product in the future? Or do they just take it at that point in time and leave it? That's critical because you can get an insight into what it really means when a vendor says they can add features.

Migo: This is a really interesting point about the maturity of the product. I've been involved with products that were more mature and some that were less mature, and what I saw is that some customers prefer less mature products because they can influence the product development and change it so that it would be suitable for them. But many security leaders are looking for tools more than solutions, in a way. They want to be able to tweak everything and they want to be able to define what to block and what not to block. And I think

there is another part of our community of enterprise customers that are really looking for a solution that, I wouldn't say is install-and-forget, but that will take most of the load from the security teams today.

It's been harder to get good security experts, simply because there's such a high demand for these people, and it's really hard to retain them. So I think we will see more solutions that are designed as solutions, rather than as a tool. A tool would be just a capability, maybe like a firewall that doesn't come with any policy, so that you need somebody who understands both the security landscape and the organizational needs start defining those rules. The solution example would be an endpoint security solution that usually comes with a predefined set of capabilities and is capable of detecting without requiring any prior knowledge from that particular enterprise.

Les: That's true. First of all, there is an influx of connected devices. Things like the Internet of Things (IoT) and mobile devices. Just the sheer amount of connected devices requires machine learning, which is here to stay and will get more mature. There are huge gaps right now and some companies are moving faster than others. We are going to have to start sharing intelligence – it's already happening in many ways. We'll have this whole orchestration of tools that have these standards – well, we call them APIs - but we need an easier way of sharing and making things better because the bad guys are doing the same thing. We have to just keep catching up and we have to be adaptive.

Migo: I definitely agree and I will be harsher than you are. In my opinion, a product that cannot integrate with other products on the same network should not be there. Products that are isolated, that do not share information, that cannot, for example, send a command to the firewall to block something and to keep the integrity of the network, a product that does not do that is really behind. If I was a buyer, I would definitely not go and buy something that cannot enrich my existing security stack.

And of course, orchestration. We've seen in the last year a few acquisitions related to orchestration and I think we will see more in that direction because

there should be a fair amount of automation within our space. Alert fatigue is something that is always an issue. Dwell time, also, is something that is always a concern. One thing, for example, that I read about the Target attack, back in 2013, is that it took 90 days for them to notice. There was actually an alert but they had so many alerts that they didn't notice for 90 days. If we don't implement these orchestration layers, this automation, I think we will be losing the battle of securing the enterprise. This is my personal opinion.

Les: Absolutely and actually Target was just one company that was exposed and in the news. But for every one firm that gets exposed, there are hundreds of others. I've seen that in my life, working at other companies. We see these alerts, and there are just so many, that it becomes a mundane thing and you stop thinking. There's an old adage if you recall it: a fool with a tool is still a fool. You need to have the skillset, maturity in processes, how you escalate things, how you pass on information to people - and the tool itself. Only then you can be successful.

You might have the best tool, which is a good thing, but you should also have the intelligence to sometimes validate. Like when something is said to be a threat. You need to have expertise. I love the 'trust but verify' mantra. But how do you do that? You have to balance the two. And for that, I think you need artificial intelligence. Moving forward, the machines will start learning stuff that we tend to do.

Migo: We already see that the bad guys are using machine learning, so it's become a commodity or a necessity for security solutions. With that said, I've seen many people focusing for example on what algorithm the AI uses. To me as a vendor, that's the wrong question to ask, even though there are a few vendors out there who lead their messaging with that. To me, AI is a tool that is dependent on the data you're using. If you have partial data or data that is not reliable enough, it doesn't matter which algorithm you're going to run it - it would not be accurate. It will not give you the results. So, I challenge my vendor peers to talk about what is the benefit of the customer rather than what algorithm we implemented. I'm not married to any algorithm. I will change my

algorithm if it would give better results for my customers. The reason we exist as vendors is to provide value to our customers. This is my view on AI.

Les: And to expand upon what you said, it's important to get that data and curate it so that the AI is learning from meaningful data. You have to consider data mining and extracting the data that is important so that the patterns can be studied and understood. Back in the day, it would be just one process or a few processes, now they are all inter-related and data is going laterally, a pattern happening here is talking to another pattern and another machine with some processes. The differentiators are that the companies may not just collect that data but make use of it with data mining and even sharing intelligence.

Migo: I agree. I think there are a few steps to what you just said. First is that taking the data and putting it in context is the first challenge because data without context is just occupying your disk space. It doesn't help anyone. And second, when a vendor has a lot of data and context, then they need to share their expertise. Without the research guys that know exactly how attackers are executing their attacks, even if I get the data of the entire planet, I would not be able to have a good product. So it's a mix between data, the context of these data and the human scale, in order to build something that is beneficial.

Les: I concur. One of the things we spoke about earlier was that there's never going to be total prevention. And that's why I think we have to be better, as a society and as companies, at asking how we manage a crisis after it happens. How do we contain whatever it is? You have to be prepared for that, otherwise, you'll be caught out. That's a big thing in the world today, whether we talk about privacy or anything of that sort. Companies that survive actually learn how to manage better after that crisis happens.

Migo: So as the leading enterprise, you have in place a kind of playbook to respond to any issues that may occur from a cybersecurity perspective?

Les: Well, I can't say exactly what we do in our company for all that. I believe some companies are better at that. Actually, the process-oriented companies, for me, are better at that than less process-oriented companies, but I think that mindset has to change. Because not all companies think that way. Most companies I've seen just react to a certain issue. From the world's perspective, it is usually the first people who get targeted who are the ones who are exposed, because somebody tries to make them the scapegoat or example. Other times it's a political reason, right? A politician wants to say, 'hey I'm taking care of my people so therefore I do these things'.

Migo: To me, that represents the maturity of an organization, to have this process assuming the worse and understanding what every function within the enterprise needs to do if something bad happens. It's a level of maturity that I really respect because I've seen companies that don't want to speak about 'what if?'. Every month we see more companies being associated with security breaches and we've even seen some cities in the U.S. affected by ransomware.

Also, the latest ransomware, LockerGoga, in Norway, and you saw over there how they really managed the situation carefully. They immediately initiated a press release, they reported on their findings. They kind of managed their brand reputation aspects. I was really impressed. I don't think you can improvise this. I think this is about organizations having a playbook or maybe even practicing. To me, it really represents the maturity of an organization that they are aware that cybersecurity is something that is here to stay. You need to keep on investing in your security posture inside the organization. To me, it's very clear that a company that does that is really doing better than others.

Les: Absolutely. And you mentioned the playbook because it's critical. OK, so I got hit. I should sometimes not say that right away, to be fully transparent. Sometimes it takes time to do a full assessment. Just because I got breached it doesn't mean that it's a problem because it could be that it's contained and that nobody else knows. Maybe we managed to contain it. A public company maybe there are legal reasons you might have to disclose a breach but you're absolutely right. It's a maturity of the process and the training behind it.

Migo: Ok, so poking your mind a little bit more, what is missing from your experience when enterprises evaluate security products? Is it the process that they don't have? Is it that they trust a recommendation from a friend, or they don't have the means to evaluate? Maybe the malware lab or the personnel?

Les: I think it's a combination of all these. Depending on how the culture is, for the most part, companies go with their past experiences. This particular vendor bailed me out in this previous job and therefore I have this close affinity, versus the method that we spoke about. Your best approach is to validate or make a case by taking your own requirements and see if they match up to what the vendor is offering. You can test the basic stuff that you want. Match it against the kind of experience and culture of your organization. Also consider trusted references. You can take something and then it's useless because you can't use it or is just, like you said, lying on the shelf. I've seen that first-hand in many companies, where they buy the product and it's just lying there.

Migo: Is there a feature that you would have expected to see by 2019? Something that enterprises should be able to purchase and use but there's no offering of it yet? Is there a gap that's not being filled?

Les: I spoke about integration and orchestration right? It is a gap today because the way orchestration is working it's just creating this one spot. But some tools today are so good at doing that first level. I wish it would happen on the whole orchestration.

The other thing I see, which is a huge gap, is the integration of mobile and perhaps using a blockchain as best we can. But there's also the whole aspect of privacy. I have seen studies (for example, such as one touted by Mastercard[5]) wherein they are creating this concept of letting the data reside wherever it is and controlling user access with a private identity token stored in their

5 https://www.mastercard.us/content/dam/mccom/en-us/issuers/digital-identity/digital-identity-restoring-trust-in-a-digital-world-final-share-corrected.pdf

possession (for example, a smartphone). Notably, you can enrich the data by feeding in additional points of identification. And the only thing that will happen is it'll be verified using elements of the cryptographic concept known as a "zero-knowledge proof," - the data never moves. If a rental company needs to know something about me, they only ask that identity. That gives you the advantage of not moving data. So it remains a better way but obviously has to be protected. That's another access control issue but that's one aspect that I'm thinking about. I was fascinated by that idea.

Lastly, social engineering and human aspect gaps are huge. There has been some improvement in tools to address some behavioral issues. But the context is much bigger because one has to take into context various environment, country, and business cultures. I think that we will continue to see increasing use of AI to learn and act upon rules for social engineering gaps. These rules could be adapted from normal behaviors, classifications, etc. I would be tempted to include quantum physics - but that is best left for another discussion that is likly to pivot AI/thinking/computation more closely to human thinking.

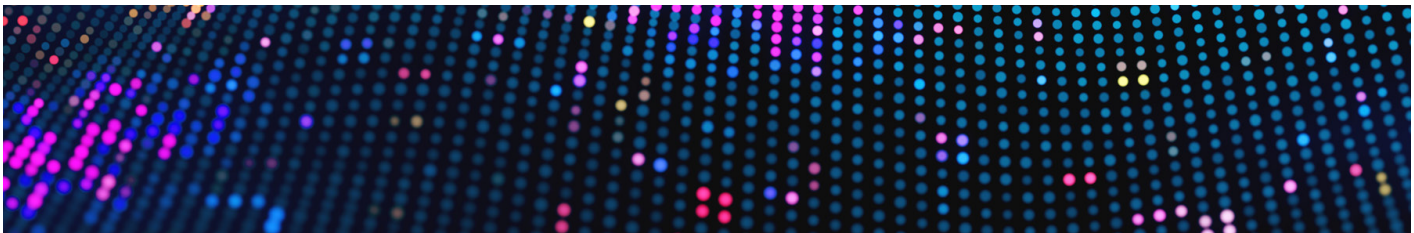

Section 7

Conclusion

The challenges of evaluating security solutions are real - while the threat landscape is evolving, introducing new zero-days and techniques to compromise devices, there are many new solutions emerging in the market, promising the earth and the moon. The CISOs, and other decision-makers who need to bridge the inherent conflicts between the business and security must also be on top of their game, mastering the risks, understanding the potential impact on the enterprise.

How do you separate the wheat from the chaff? In this ebook, we don't just talk about the problem from both vendor and enterprise client perspectives, but we also offer easy to integrate tools that you can immediately implement into your evaluation process. We covered the need to define your success criteria as early as possible in the process, which can set you in the right direction, the pitfalls, how to test and evaluate and more.

With that said - every organization has its own pain points and sensitivities - without understanding these, forming a coalition which allows more decision-makers to be part of the process, you are taking the risk of backfire.

A CISO's work is never done. That means a certain amount of your time will always be spent on evaluating new products. We hope that this book has given you some strategies for doing that so that you can choose the best tools for your needs and do not waste valuable time. We wish you every success!

Appendix A

Table - Capability/ Feature, Priority, Ranking

Priority to organization – (H)igh = 2, (M)edium = 1, (L)ow = 0

Score Ranking: 5 – Excellent, 4 – Good, 3 – Average, 2- Fair, 1 – Poor, 0 – Non-existent

Research/Evaluated – implies whether researched or Proof-of-Context test

Sum up Priority and Score Ranking to determine final ranking.

N°	Capabilities/Features	Explanation	Priority H=2/ M=1/L=0 H/M/L	Product	Score 5/4/3/ 2/1/0	Research/ Evaluated	Comments
1	Security Effectiveness	The effectiveness of detecting and preventing					
1.1	Capability to protect against threats inbound before execution, during execution and post-execution	Inbound threat detection and prevention, Execution-based threat detection and prevention, credential theft prevention, Continuous monitoring post-infection and ability to act in the event of a compromise					

N°	Capabilities/Features	Explanation	Priority H=2/ M=1/L=0 H/M/L	Product	Score 5/4/3/ 2/1/0	Research/ Evaluated	Comments
1.2	Validate each type of threat category against associate threat vectors	Threat categories: Malware, Exploits, Blended threats for the following threat vectors: HTTP, HTTPS, P2P, Email, Local Intelligence, Blended threats, Evasion, Connected IoT and Devices (USB), lateral movement					
1.3	False Positives	How effective against known legitimate traffic					
2	Detect and Response (EDR)						
2.1	Event Reporting capabilities	Support for an incident response such as Ranking of threats, API Calls, Data exfiltration, File system, Lateral movement, Registry changes, system integrity, patient 0, forensics, etc.					
2.2	Automation capabilities	First level SOC automation - ability to identify malicious acts in real time, automating the required responses and allowing seamless threat hunting by searching on a single IOC.					
2.3	Remediation capabilities	Support for host isolation, process termination, file and process isolation, repair, recovery, rollback, etc.					

N°	Capabilities/Features	Explanation	Priority H=2/ M=1/L=0 H/M/L	Product	Score 5/4/3/ 2/1/0	Research/ Evaluated	Comments
2.4	Threat Hunting capabilities	Search for cyber threats or validate reports, IOC search, encrypted traffic visibility, etc.					
3	Intelligence feeds utilized	Number and reputation of intelligence feeds					
4	Auditing and Compliance						
4.1	Complies with Regulatory Requirement for the organization	Vendor meets all regulations such as SSAE 16 SOC 2 the or ISAE 3402 Type 2 along with those required downstream by an organization in support of PCI, SOX, FDA, GDPR, Data retention, localization, Disaster Recovery, Encryption, etc.					
4.2	Audit Logging and retention	Ability to retain logs					
4.3	File Integrity Monitoring	Ability to report file changes and report					
4.4	File-based scan	Required by some regulations and standards					
4.5	Application inventory and vulnerabilities	Ability to inventory applications and vulnerabilities					
4.6	Rank Vulnerabilities	Ability to rank/ prioritize remediation efforts					

N°	Capabilities/Features	Explanation	Priority H=2/ M=1/L=0 H/M/L	Product	Score 5/4/3/ 2/1/0	Research/ Evaluated	Comments
4.7	Application vulnerability research reference	Identify sources of information and comply with Common Vulnerabilities and Exposures (CVE) and others					
4.8	Patch management	Provide links or patches for remediation of application vulnerabilities					
5	Operational controls						
5.1	Device encryption orchestration	Ability to orchestrate BitLocker Device Encryption, FileVault full-disk encryption, and others					
5.2	Firewall controls	Ability to control endpoint firewalls and manage in/out network					
5.3	Device Control	Ability to control attached devices					
5.4	Endpoint Remote Shell access	Ability to provide administrators full encrypted logged access to an endpoint command line					
5.5	IoT discovery and control	Ability to discover IoT devices in the vicinity and indicators of threats					
5.6	Data Loss Prevention	Ability based on classification parameters, ensure that sensitive data is not lost, misused, or accessed by unauthorized users					
5.7	Group management	Ability to group endpoints					

N°	Capabilities/Features	Explanation	Priority H=2/ M=1/L=0 H/M/L	Product	Score 5/4/3/ 2/1/0	Research/ Evaluated	Comments
5.8	Integration with Directory services	Active Directory and other directory services to utilize dynamic groupings/objects defined/changed					
5.9	Granular policies enforcement based on administrator selection	Ability to control policies based on organizational requirements include inheritance etc.					
5.10	Role Based Access Control	Ability to provide console access based on roles and responsibilities					
5.11	Delivery platform flexibility – Cloud, On-premise, and hybrid	Console and data collection points - multiple points increases versatility					
	Management Console	Ease of use, intuitive - Usability & Customization ability					
5.12	Mobile Admin app	Ability to enable management of the console on a mobile device					
5.13	Management Report Generation	Standard reports generated are detailed and are they customizable?					
5.14	Scalability and Growth	Ability to create multiple sites and models					
5.15	Endpoint Platform(s) Coverage	Microsoft Windows, Linux OS families, MAC OSX, Embedded systems, Cloud PAAS, VDI, etc.					

N°	Capabilities/Features	Explanation	Priority H=2/ M=1/L=0 H/M/L	Product	Score 5/4/3/ 2/1/0	Research/ Evaluated	Comments
5.16	Integration capabilities via API/SDK	Ease of integration and current integration with tools/ platforms compatible with the environment to share intelligence and dashboard - mobile, CASB, SIEM, Threat intelligence, IT/Ops mgmt, Firewalls, Orchestration, etc.					
6	Agent characteristics						
6.1	Malware Sandboxing location	Local or cloud-based – do you need connectivity to the cloud/ platform?					
6.2	Agent Weight	Size of agent					
6.3	Agent to Cloud/console bandwidth consumption	Traffic consumption – agent to cloud/platform					
6.4	Agent scalability	Multiple agents or a single agent to perform capabilities?					
6.5	Agent Configuration and Update	Ability to control agent memory, capability, local space usage, etc.					
6.6	Ease of Deployment	Ease of deploying agent and vendor support (gratis/PS services) – third-party tools, GPO, scripts, etc.					
6.7	Reboot for the first installation?	Does the first installation require a reboot?– and services enabled with a reboot?					

N°	Capabilities/Features	Explanation	Priority H=2/ M=1/L=0 H/M/L	Product	Score 5/4/3/ 2/1/0	Research/ Evaluated	Comments
6.8	Reboot required for Updates?	Do updates require a reboot?					
6.9	Online/Offline efficacy	Efficacy if online/offline (disconnected)					
7	Vendor background						
7.1	Credibility and stability	Years in business, success stories, global reach, stability, responsiveness to changes/issues, change management validation (promise versus delivery)					
7.2	Support Organization	Responsiveness, communications, support options, and SLAs					
7.3.	MDR Services	Ability to provide services that combine threat intelligence, endpoint/network data, security hygiene, and anomaly information					
8	Financials for investment	Costs, savings, Operational costs, CAPEX/OPEX and ROI – estimate 1-3 year					
8.1	Costs	Estimated license purchase, maintenance, Threat alerting and monitoring, resources					
8.2	Savings	Total savings because of product (EPP, EDR, and/or Incidence Response). Consider whether this product will replace multiple products.					

N°	Capabilities/Features	Explanation	Priority H=2/ M=1/L=0 H/M/L	Product	Score 5/4/3/ 2/1/0	Research/ Evaluated	Comments
8.3	Operational expenses	Costs to be allocated to decrease the likelihood of compromise					
8.4	Return on Investment (ROI)	Calculate based on costs, savings, and operational expenses					
9	Research						
9.1	NSS Labs rating					R	
9.2	Gartner Report rating					R	
9.3	Forrester Wave Report rating					R	
9.4	IDC report					R	
9.5	Av-test.org report					R	
9.6	www.amtso.org - reports					R	
9.7	False positives report and efficacy	Look up sites such as VirusTotal				R	
9.8	Vendor presentations and demos	Insight into product				R	
9.9	Research and networking intelligence	Any legal or known customer feedback from network, news etc.				R	

N°	Capabilities/Features	Explanation	Priority H=2/ M=1/L=0 H/M/L	Product	Score 5/4/3/ 2/1/0	Research/ Evaluated	Comments
10	Product Road Map - Determine whether the vendor's growth path for the product aligns with organizational needs						
11	Organization specific considerations	Incumbent vendor, current environment complexities, contracts, dependencies of products, coexistence requirements etc.				R	
12	Vendor References	Sample questions and responses of note: How long using the solution, previous vendor, number of endpoints, operational issues resolution and support experience, coverage, interdependencies, coexistence with other products, number of organization resources supporting, roadmap timeline validation, deployment issues, lessons learned, and timeline.				R	
	Final Ranking = Sum of Priority + Score Ranking						

Appendix B

Malware Testing

This section provides a short introduction to testing endpoint solutions. It focuses on attacks post-delivery and not necessarily on attack vectors, such as social engineering. It is critical to realize that the more features an endpoint solution has to prevent, detect, and remediate aspects across the attack chain, the more effective it will be in defeating the threat. This method may be lost if you do not understand the context of infection propagation. For example, how did it reach the endpoint? Additionally, with the increase in fileless malware, this lends context that is critical.

Proper malware testing with unknowns or 0days must follow testing criteria. Testing unknown, or malware that is new to the community (0days) is the most prominent blind spot from a critical security perspective for properly testing malware. There are many resources for pulling down malware samples that you can then use for testing after you have mutated them to make 0days. Check the following sites for malware, documentation, and methods:

http://Virustotal.com - registration required and a paid subscription for multiple (e.g., top 50, etc.)
https://www.shellterproject.com/introducing-shellter/
http://Offensivecomputing.net
https://github.com/ytisf/theZoo
https://github.com/ytisf/theZoo/tree/master/malwares/Binaries
http://malwaredb.malekal.com
https://testmyav.com/
https://virusshare.com/
https://www.phishtank.com/
http://tracker.h3x.eu/
https://malpedia.caad.fkie.fraunhofer.de/
http://vxvault.net/ViriList.php

http://malc0de.com/database/
http://support.clean-mx.com/clean-mx/viruses.php
http://virussign.com
http://malwaretips.com
http://contagiodump.blogspot.com/
https://github.com/curi0usJack/luckystrike
https://kakkulearning.wordpress.com/2014/07/06/add-new-exploits-metasploit-exploit-db/
https://malshare.com/
https://attack.mitre.org/wiki/Main_Page

Note:

- Downloaded malware must be stored in a directory that does not have execute rights. Ensure that this directory is in the exclusion list of the product(s) being tested.
- If moving malware between machines, ensure that they are zipped and password protected.

Bring Your Own Malware (BYOM) Guidelines

It is important to use fresh malware samples as opposed to those that have been in circulation for some time (with over 500,000 new malware samples generated per week, according to AVTest.org, this should not be a challenge). Ensure that the malware sample executes – there are well-known instances of so-called "fresh" malware samples that are corrupted files that will not run or are malware that is rendered benign because the Command and Control (C2C) server have been taken down.

If you are designing your own malware, make sure it resembles a real attack. For example, assemble a PDF file, sent via email, have the end-user open the PDF, that makes an outbound connection to a web-page that compromises

RAM, injects code, or streams commands to the shell, establishes persistence (maybe something simple as modifying *autorun*) and then tries to delete, or siphon a file from the system to a remote location.

There are other options. You can make your own malware mutations via many tools that maintain the payload or tweak them to make a new 0day - such as:

1. Using packers and similar malware mutation methods:

    a. **AEGISCrypter** – see *http://www.aegiscrypter.com/*

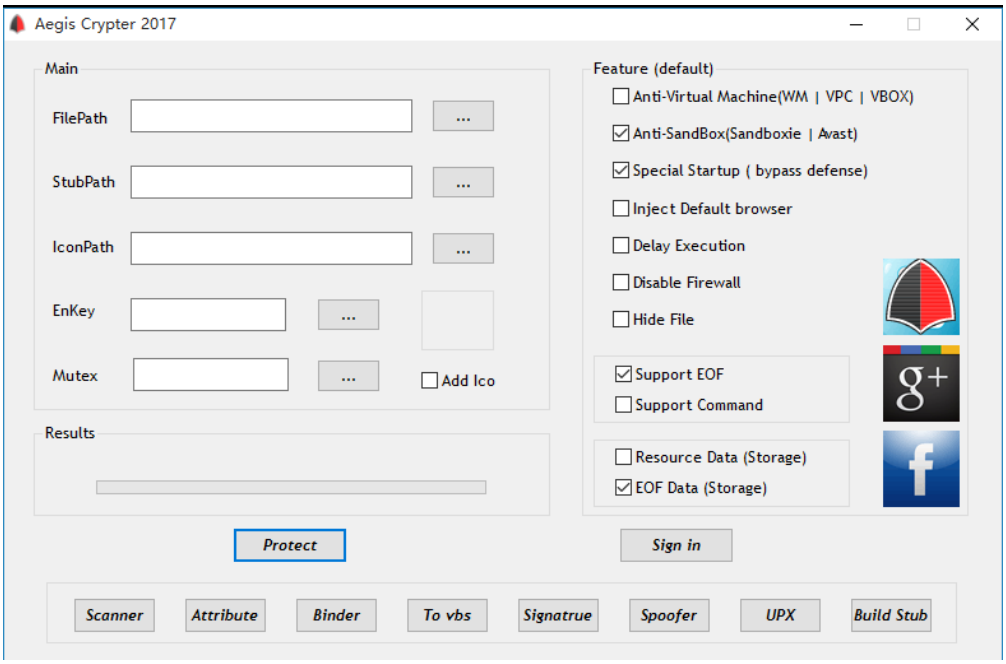

Figure 6 - Aegis Cypter

Instructions on usage are self-evident on the website and screen sample.

b. **MPRESS** – see *http://www.matcode.com/mpress.htm*

The official site is sometimes down. Google for other sources if it is *https://www.google.com/webhp?sourceid=chrome-instant&ion=1&espv=2&ie=UTF-8#q=mpress+packer+msassist*

Instructions are available on the website or follow the instructions below:

Take any malware sample and using MPRESS create mutations on the fly. The payload will remain exactly the same as it was before, but the hash value will change thus making it unknown or an 0day.

i. Create a copy of the folder containing your malware samples.
ii. Copy mpress.exe to that folder.
iii. Open a command prompt and change directory to the new malware folder.
iv. Run the following command: for %i in (*) do mpress –i %i
v. When it is done packing all the samples, open the folder and sort by file modification time so old files are on top. You should be able to delete mpress.exe from the folder and any samples that failed to pack by deleting the older files.
vi. Drag and drop samples to each machine.

c. **VMPprotect** – see *http://vmpsoft.com/products/vmprotect/*

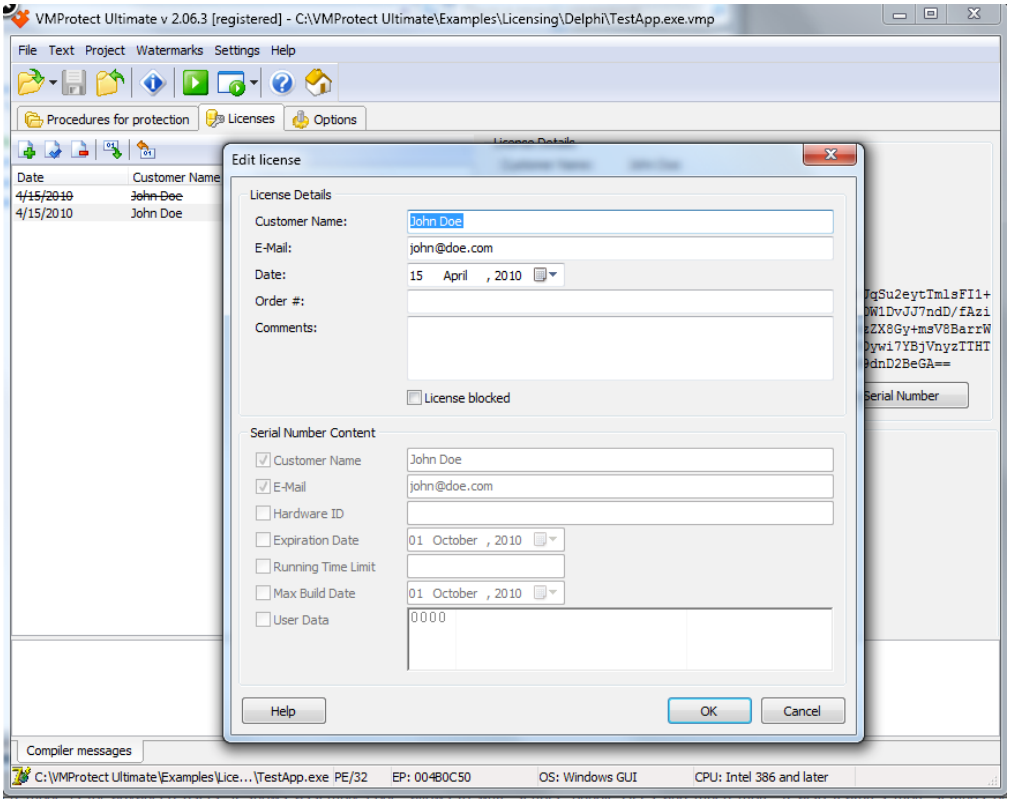

Figure 7 - VMProtect

Instructions on usage are self-evident on the website and screen sample.

2. Simulate "unknown" malware by installing the product solution agent, and then unplugging the network cable, wait for a few days, download new malware from the malware websites to USB, and then run tests.

3. Utilize Kali Linux to test some exploits. See *https://www.offensive-security.com/metasploit-unleashed/*

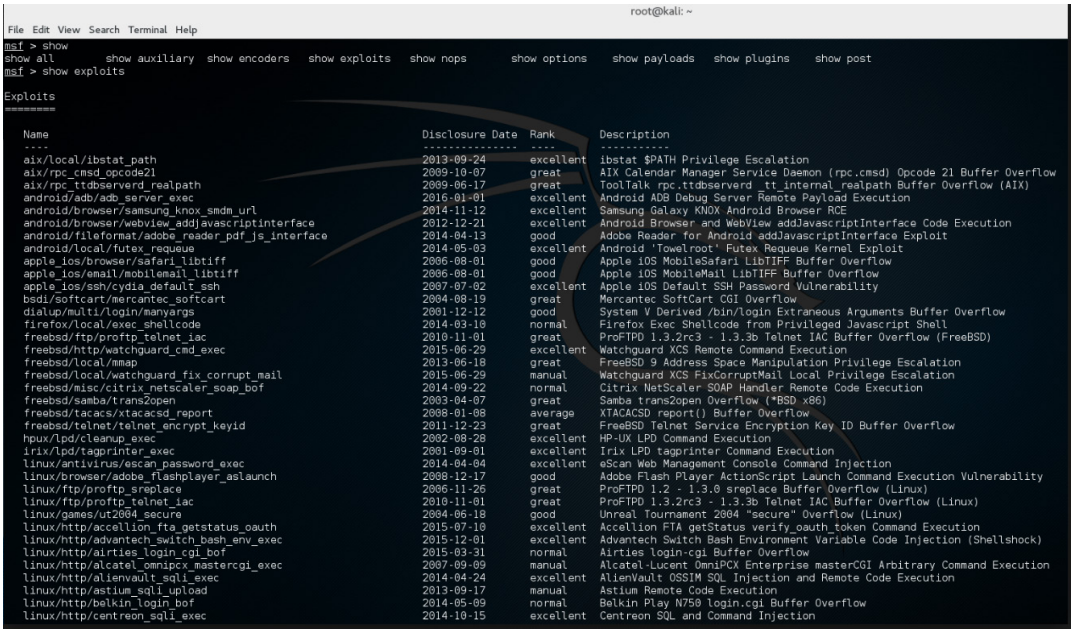

Figure 8 - Kali Linux exploits sampling

The typical testing process might look like this – assuming testing on a victim VM with the EPP/EDR installed in the isolated lab:

1. Test each exploit against the victim VM, and rolling back to the "clean" snapshot.

2. Assess each exploit attempt:

    a. Did the exploit succeed?

    b. Did the EPP/EDR software raise an alert within to the?

    c. Did the EPP/EDR raise an alert within the admin console?

        i. What triggered the detection? Was it the Metasploit and what specific exploit technique or signature?
        ii. Did the EPP/EDR take to block or disrupt the exploit?
        iii. Did the EPP/EDR take to remediate?

d. Did the exploit fail but without any alert from the EPP/EDR product?

   iv. Are you sure the exploit works reliably without the endpoint protection software installed?
   v. Does the EPP/EDR software enforce preventive exploit mitigations, such as Address space layout randomization (ASLR), without triggering an alert?

There are a number of techniques that are explained in the links shared.

4. Get malware/research techniques from your Security researcher friends


SentinelOne™