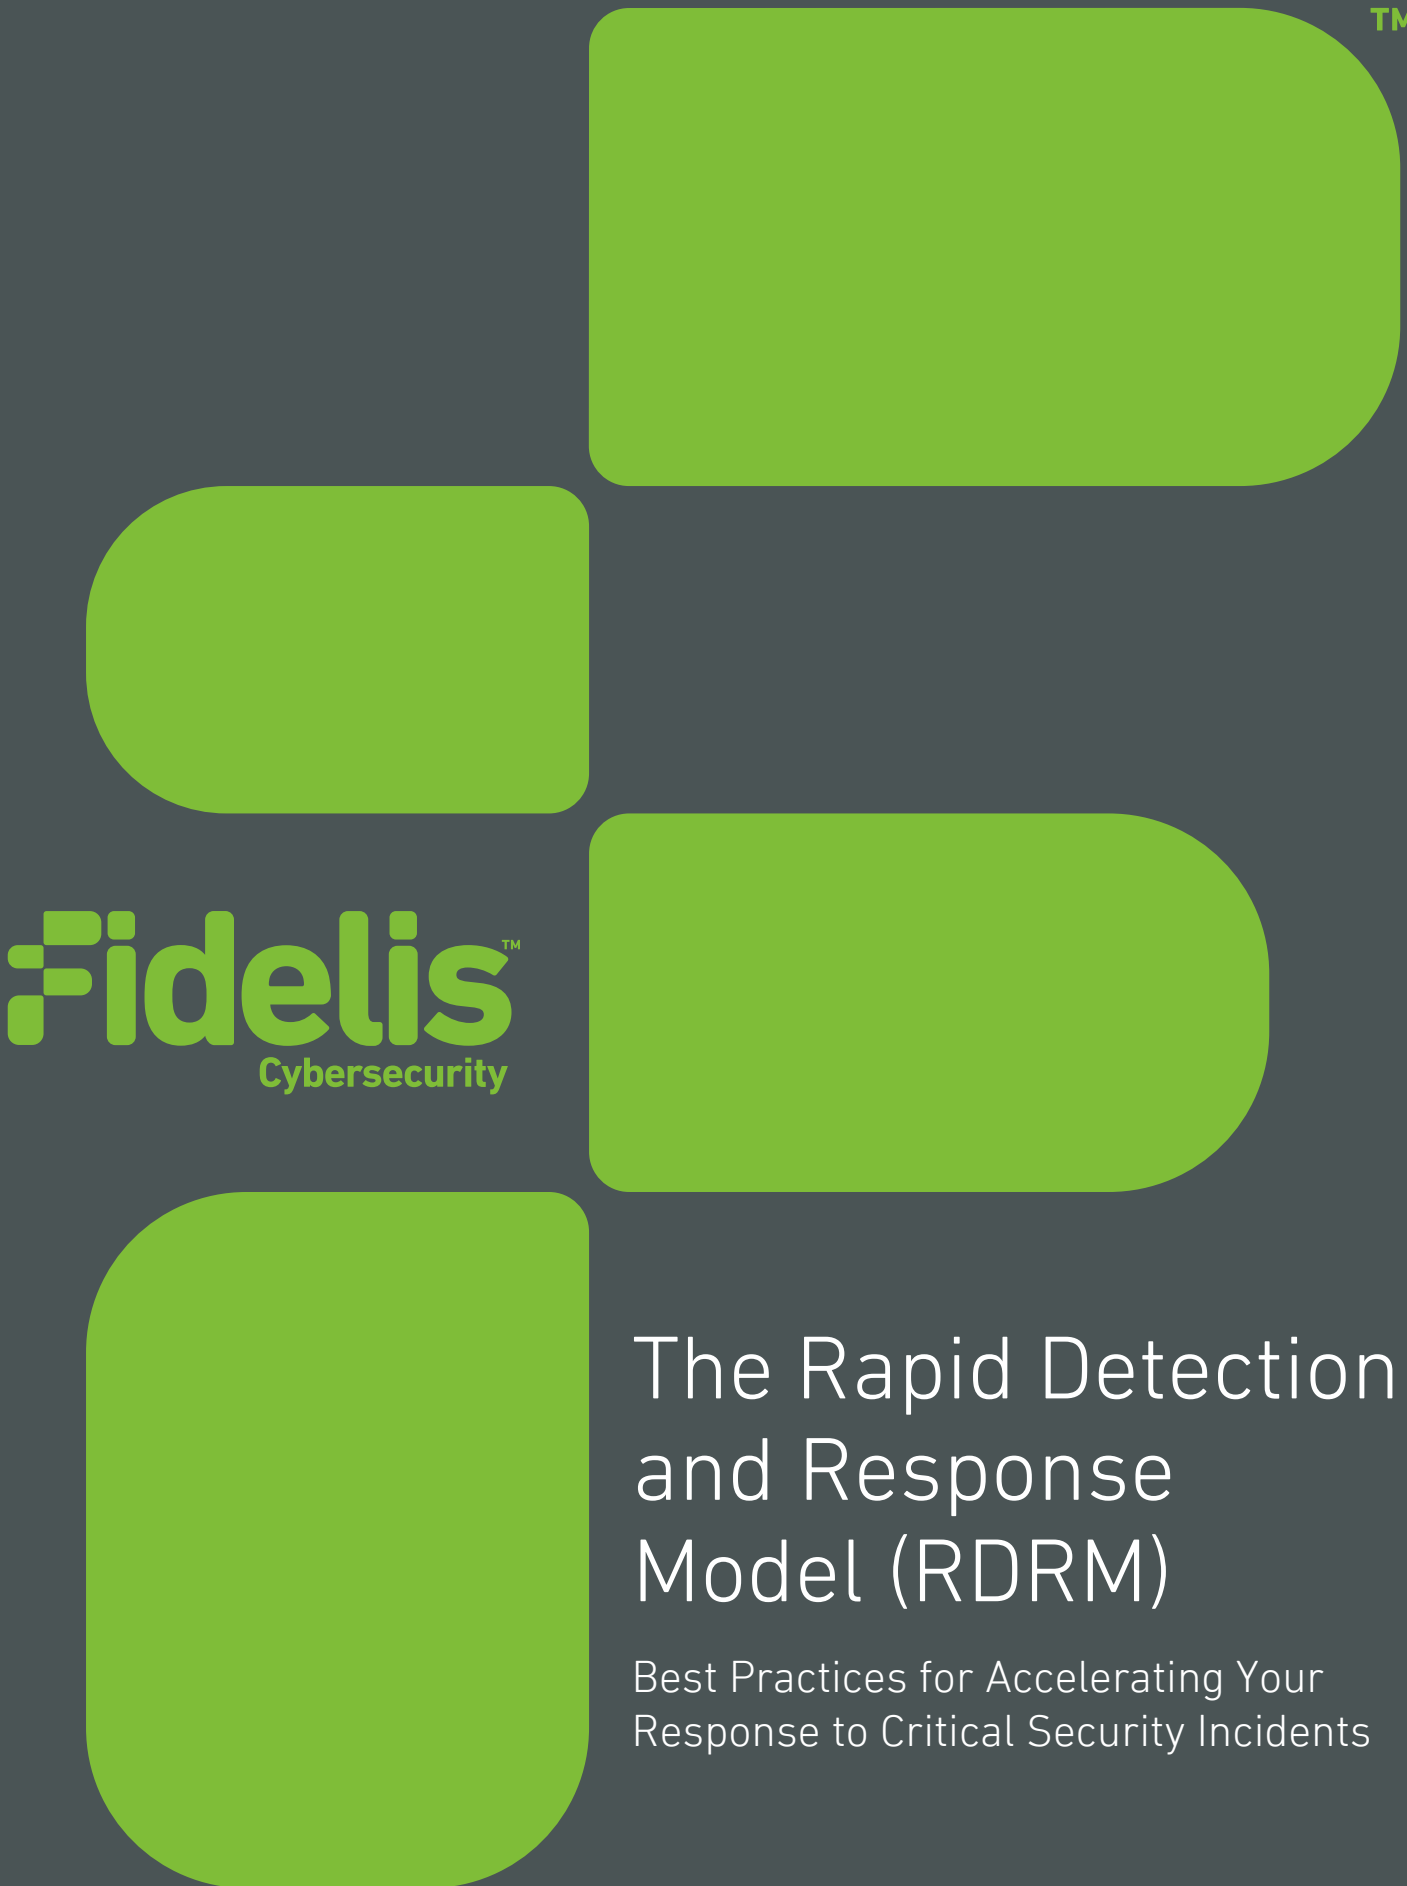

The Rapid Detection and Response Model (RDRM)

Best Practices for Accelerating Your Response to Critical Security Incidents

www.fidelissecurity.com

Contents

Executive Summary

The number of threats organisations face is increasing exponentially as attackers — armed with new tools and techniques and inspired by a range of motives — grow more sophisticated in their actions. Despite investments to build secure networks, determined attackers continue to routinely compromise seemingly secure organisations and steal their intellectual property, private data and financial information.

Most security teams lack the manpower, visibility and threat intelligence needed to quickly and accurately investigate the huge volume of alerts generated by existing security solutions. The ability to proactively hunt for the advanced attackers that silently manoeuver past defences is generally on the wish list. As a result, most security breaches go undetected until it's too late, and an outside incident response team is often brought in to assess and remediate the damage.

This guide details the Rapid Detection and Response Model (RDRM). It shows how organisations can overcome these challenges so they can more effectively detect and respond to network intrusions and data breaches. The model promotes more efficient incident detection and response by advocating for the consolidation and integration of endpoint, log file and network visibility technologies. By adopting the RDRM and implementing the recommended practices, organisations can measurably reduce their risk by aligning their systems and processes to accelerate their ability to detect, investigate and resolve critical security incidents.

By adopting the Rapid Detection and Response Model and implementing the recommended practices organisations can measurably reduce their risk by aligning their systems and processes to accelerate their ability to detect, investigate and resolve critical security incidents.

Threat Detection and Response — a Daunting Task

Security analysts and incident responders, overwhelmed by alerts and tasked with reviewing and prioritising suspected incidents, are challenged in their ability to quickly detect threats as they are happening and validate whether a suspected incident is real or not. Further, they receive little context on the potential impact. Often, network security solutions are not linked to endpoints — leaving security teams blind to all but the most basic information about whether desktops, servers, laptops or mobile devices across the organisation have been compromised or if attackers are using them as a launch point. This makes it challenging to join the dots or know where to allocate limited resources in order to respond appropriately to a suspected incident.

When analysts do respond to an incident, the task of determining which systems are potentially compromised and retrieving the data can take days. Performing manual investigations is time-consuming and typically requires using multiple point solutions to view network traffic, endpoint data, threat intelligence and other data sources and then reconstruct what happened. With staffing stretched thin, there are simply not enough qualified resources available to adequately keep pace with attacks as they increase in volume and complexity.

The net result is that analysts often miss the most critical attacks or detect them long after vital data has been stolen. Why? Signs of an initial attack can be stealthy and are difficult to differentiate from the noise of the day-to-day deluge of alerts.

Combating today's cyber threats requires an intelligence-driven defence and technology that equips analysts with the context, visibility and speed required to confidently identify, investigate and stop advanced attacks.

The sheer number of genuinely important alerts makes it nearly impossible to respond to all of them and manual triage processes slow teams down. Delayed response times and inaccurately prioritised alerts further compound the problem and create gaps that attackers use to gain a foothold and roam freely across a network. As time passes, the cost of an incident spirals upward as victim zero leads to victim one hundred, dozens of backdoors are planted, account passwords are stolen and data theft occurs.

Combating today's cyber threats requires an intelligence-driven defence and technology that equips analysts with the context, visibility and speed required to confidently identify, investigate and stop advanced attacks. By automating incident response and using technology to improve the efficiency and accuracy of security analysts and incident responders, organisations can reduce the time necessary to detect and resolve incidents so they can more effectively mitigate risk.

The Rapid Detection and Resolution Model (RDRM)

The RDRM is designed to lower the risk profile of an organisation and increase efficiency with measurable results by ensuring the organisation is prepared from a people, process and technology perspective.

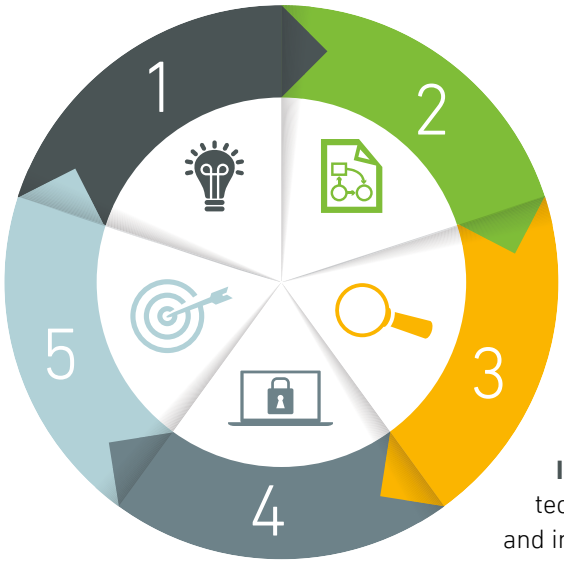

To accomplish this, the model comprises five steps designed in a feedback loop:

Identify: Creates situational awareness by identifying technology and process gaps that lead to blind spots and inefficiencies. Here, security teams document existing security infrastructure, analyse the capabilities of security technologies, examine operational processes, review detection and response metrics and evaluate the threat landscape.

Prepare: Closes gaps critical to rapid detection and response by implementing technology, integrating systems, modifying processes and performing tabletop exercises to train personnel.

Detect: Identifies security incidents. At this stage, security teams monitor and apply threat intelligence to endpoints, network traffic and log files to validate alerts; and perform security analytics to uncover suspicious anomalies.

Respond: Entails investigating security incidents to understand what happened. Here, investigators contain affected systems, collect and analyse data to classify the threat, dissect the attack path and reconstruct what happened.

Resolve: Creates and implements a remediation plan to remove all points of entry available to the threat. In this phase, responders remove backdoors, fix exploited vulnerabilities and reset compromised user credentials. Teams should also document lessons learned and new threat intelligence to refine and improve the model.

Visibility, Integration and Automation

With cyber threats growing exponentially, security teams are challenged to keep pace. The traditional information security infrastructure is typically a fragmented patchwork of network, endpoint, security information and event management (SIEM) tools and other security infrastructure that often fail to interoperate well. To provide security teams with the ability to rapidly and continuously detect, quarantine, respond to and resolve incidents, we recommend aggressively pursuing three technology goals:

- **Visibility:** Gaining actionable visibility into logs, networks and endpoints:
  - Increases insight across the enterprise for detection and response purposes
  - Provides context to help validate security alerts and understand incidents
  - Decreases risk by making incident detection and response more accessible
- **Integration:** Using a single platform with multiple capabilities that interoperate and can apply threat intelligence to your environment:
  - Decreases the required number of endpoint agents
  - Reduces the number of manual steps required to piece together data from multiple sources
  - Facilitates correlation and review of relevant log, endpoint and network data
  - Provides "big picture" context that can accelerate and improve threat detection accuracy
  - Enables automation to span multiple products
  - Stretches budgets by reducing the number of niche point products
- **Automation:** Automating tasks and workflows frees experienced security personnel to focus on high-priority tasks:
  - Validates alerts
  - Contains compromised endpoints
  - Collects data from endpoints, network traffic and other relevant sources
  - Performs preliminary analysis of both data and malware
  - Provides analysts with the data and context necessary to prioritise and take appropriate action
  - Pushes remediation actions to affected systems

Step 1: Identify

KEY WORKSHEETS

Appendix A: Threat Activity

Appendix B: Rapid Detection and Response Metrics Self-Assessment

Appendix C: Rapid Detection and Response Model

The purpose of the Identify step is to create situational awareness of the organisation's threat environment. It establishes a baseline understanding of your ability to manage cybersecurity risks and your organisation's incident response maturity level. The Identify step includes and extends requirements from the Identify and Detect functions of the U.S. National Institute of Standards and Technology (NIST) Cybersecurity Framework with a focus on what's needed for rapid threat detection and response.

Analyse Threat Activity

We encourage you to begin by analysing the nature of the threat activity in order to understand its severity and identify trends. *Appendix A, Threat Activity*, provides a worksheet that you can use to record the actions taken by threat actors as their attacks progress through your network.

Assess Detection and Response Metrics

Assess your ability to rapidly detect and respond to threats. *Appendix B, Rapid Detection and Response Metrics Self-Assessment*, provides a worksheet to record your answers. Be sure to add to the questions below with some that are unique to your organisation.

- How long does it typically take to:
  - Discover threats that bypass defences?
  - Validate a security alert?
  - Contain affected endpoints?
  - Collect the endpoint, network, logfile and application data needed for analysis?
  - Analyse and identify the threat, root cause and scope of the incident?
  - Resolve the incident?
- What percentage of incidents goes undetected until obvious symptoms appear?
- What percentage of incidents is thoroughly investigated to answer important questions?
- Do these numbers change for remote workers and satellite offices?

Evaluate Security Processes and Technologies

Finish by creating an inventory and evaluating the operational capabilities of your security processes and technologies so you can identify gaps that lead to blind spots and opportunities for improved efficiency. *Appendix C, Rapid Detection and Response Model*, provides a worksheet to document your current capability and prioritise desired capabilities in the following categories:

- Processes
- Endpoint visibility
- Network visibility
- Log management
- Threat intelligence

Step 2: Prepare

The Preparation step makes use of the analysis and situational awareness obtained in the Identify step to close gaps that hinder your ability to efficiently detect, respond to and resolve incidents.

Many organisations have invested in a collection of security technologies, but may not be experiencing the full benefit of their investment due to poor integration, unnecessarily complex processes or unused functionality. Also, organisations often put security tools in place as a reaction to a breach instead of in preparation for one. The RDRM helps you accelerate rapid detection and response by focusing attention on technology that makes security personnel better and faster. The most crucial features are the ones that drive down the time, cost and manpower required to detect an incident, collect relevant data from multiple sources, analyse the data, classify the threat, scope what was compromised and resolve the incident with a high degree of confidence.

Where to begin

A basic rule of thumb of the RDRM is to first focus on reducing noise from existing systems and then take steps to drive down key metrics for incident response and resolution to acceptable levels (i.e., time-to-detect, time-to-validate, time-to-contain, time-to-collect, time-to-analyse and time-to-resolve). We recommend that you give priority to workflows that you can accelerate with improved visibility, integration and automation. Additionally, now is a good time to review processes for inefficiencies or delays so that you can take action to remove or modify them. For further insight combined with survey data, we recommend "Incident Response Metrics"[1] by Sean Mason and "Incident Response: How to Fight Back"[2], by Alissa Torres, a certified SANS instructor specialising in advanced computer forensics and incident response.

Close threat detection gaps

The time between when an asset is compromised and when you identify the incident is the time-to-detect. This is also known as the dwell time. The ability to reduce dwell time is arguably one of the most important aspects of rapid detection and response. However, the following three detection gaps make it challenging for organisations to detect attackers at every stage of the attack lifecycle:

- Gaps in visibility
- Lack of process, insufficient staffing
- Lack of strategic and tactical threat intelligence

1 Mason S., "InfoSec Insights: Incident Response Metrics," 2014 [online] http://seanmason.com/2014/07/14/incident-response-metrics/

2 Torres A., "Incident Response: How to Fight Back," 2014 [online] https://www.sans.org/reading-room/whitepapers/analyst/incident-response-fight-35342/

The ability to reduce dwell time is arguably one of the most important aspects of rapid detection and response.

Gaps in visibility

The most common and riskiest gap that leads to long detection times are gaps in visibility into endpoint activity, network activity, log files and applications.

Endpoint

Most endpoint security products are designed to block the execution of processes or exploitation of application vulnerabilities. They lack the ability to import machine-readable threat intelligence and use it to find attackers. They also lack the data collection and analysis capabilities that you need to determine the root cause and reconstruct an attacker's activity. To ensure adequate visibility on your endpoints, the RDRM encourages you to implement Endpoint Detection and Response (EDR) technology that:

- Validates an indicator of compromise (IOC) and automatically sweeps endpoints for signs of compromise
- Includes all operating systems in use, including mobile operating systems (Android, iOS)
- Automatically triggers an alert any time a threat indicator is found on an endpoint

Network activity

Traditional security approaches primarily focus on preventing breaches. Unfortunately, with this approach, once attackers get through the perimeter, there is no way to detect them. That is because most network security solutions only look at a few common application protocols and services, such as HTTP, SSL, FTP, DNS and SMPT. Attackers know this and once they have established a foothold, they exploit this shortcoming. Internet traffic visibility for threat detection is less common thanks to firewall and network IDS logs. Plus, a lack of intranet logging for services such as DNS and DHCP is a key contributor to lengthy incident response and resolution times as investigators struggle to map Internet traffic to internal resources. This is especially true when reconstructing attacker activity that is weeks or months old. DNS and DHCP log file rotation means that often the required data simply does not exist. To achieve full visibility into network data, implement advanced threat-detection technology that:

- Monitors common initial infiltration vectors, such as Internet traffic, email and web servers
- Inspects network traffic on all ports and protocols, including the misuse of protocols and services on non-standard ports
- Decodes and analyses content in real-time, no matter how deeply embedded
- Provides the capability to detect and investigate retroactively
- Performs analytics on netflow traffic

Log files

Lack of visibility into centralised log files using a SIEM is still a common gap. Many organisations fail to collect the right type of events or omit important sources. To improve visibility into log files:

- Consolidate and aggregate logs
- Index and store machine data from all sources, including cloud applications
- Configure logs to capture events critical to reconstructing attacker activity (netflow, authentication success/failure and process execution)
- Set real-time alerts for potential threats
- Apply analytics to correlate patterns for known and unknown threats
- Centralise and preserve critical system, network device and application logs

Applications

Application visibility is often lacking, especially with web applications. Further, the criticality and confidentiality of web applications means more organisations rely heavily on SSL encryption. As the encrypted traffic increases, IT becomes increasingly blind to the traffic — particularly SSL interactions between enterprise users and external applications. This blind spot raises many security concerns as attackers can potentially exploit SSL traffic and move within these encrypted tunnels.

To improve visibility into web applications:

- Create a global inventory of all your public-facing web applications
- Perform a comprehensive deep scan to identify exploitable vulnerabilities
- Add SSL inspection capabilities to the network security architecture to close the security visibility loophole created by encrypted traffic
- Customise web application firewall rules to identify and block attacks (such as cross-site scripting (XSS) and SQL injection)

Lack of process, insufficient staffing

Another common problem in threat detection is missed alerts or shortcuts taken for incident response. Failing to determine the root cause of alerts or reimaging systems without investigating them leaves important questions unanswered. You could miss the opportunity to identify a larger incident that involves other systems, compromised credentials and stolen data.

In cases where the incident is only partially detected, inaction is usually due to:

- A lack of process requirements to answer important questions about an alert or detected incident
- Insufficient staffing in the face of a large volume of alerts, leading to alert fatigue
- Insufficient time and expertise to identify and validate alerts

Ways to remedy a lack of process and insufficient staffing include:

- Establish or modify threat-detection processes to eliminate gaps in procedures
- Conduct training to ensure that analysts are familiar with what normal looks like
- Implement integrated platforms that automate alert validation, data collection and analysis

Lack of threat intelligence

Threat intelligence consists of knowing who the threat actors are: their motives as well as their tactics, techniques and procedures (TTPs). There are two types of threat intelligence: strategic and tactical.

- **Strategic:** This intel is based on human observations, analysis and conclusions drawn from a number of sources. These input sources could be "boots on the ground" in the field, digital media, historical data and comparative threat intelligence. It can be used to identify a threat group; its tactics, TTP's; and in some cases: give the "why" or motivation behind an attack. In rare cases, strategic threat intelligence can telegraph or predict an adversary's next move.
- **Tactical:** This intel is codified in terms a machine can interpret and apply to a given problem. There are various formats for tactical threat intelligence, but for the most part we can classify these into network-based indicators, host-based indicators and methodology-based indicators.

Intelligence is critical for prioritising defensive measures and increasing your capabilities to detect and respond to an attack in progress. TTPs encompass how an attacker generally progresses during the attack lifecycle, what they do during each phase of an attack and specific actions they take during the attack, such as the malware they use and the commands they send during interactive sessions. Defenders can use this in-depth knowledge of attackers' TTPs to their advantage. More information on threat intelligence can be found in "Intelligence-Driven Computer Network Defense," a white paper published by Lockheed Martin.[3]

To increase threat intelligence capabilities:

- Codify attacker activity and tools observed during the investigation of internal security incidents and use it to detect other attacks
- Seek out new threat intelligence sources
- Consume threat intelligence through multiple open threat intelligence standards and commercial threat feed providers
- Use service providers to monitor and identify threats in log files, network traffic and endpoints using proprietary threat intelligence
- Increase the number of threat intelligence formats that can be read and made actionable at endpoints, network traffic and centralised log files (e.g., OpenIOC, STIX, YARA)
- Make use of information-sharing groups, such as ISACs and InfraGard, for industry-specific threat intelligence

Prepare for rapid response

A good exercise that can help you prioritise your efforts is to compile a list of the types of incidents your security team is currently identifying and validating. Common examples are alerts generated by AV, network IPS/IDS, firewall and next-gen threat-detection products. A best practice when you review existing processes is to interview security personnel so you can understand the workflows they perform.

Once you understand where most of your team's time and energy is being spent you can start to close technology gaps and automate mundane and manually intensive parts of the incident handling process. Using fewer products that integrate with one another will make it easier to automate these processes. While standalone point products may offer best-of-breed features, they can also incur technical debt due to a lack of integration and require manual steps to complete workflows. Evaluate existing technologies to see if they have unused

3 Hutchins E., Cloppert M., and Amin R., "Intelligence-Driven Computer Network Defence," 2014 [online] http://www.lockheedmartin.com/content/dam/lockheed/data/corporate/documents/LM-White-Paper-Intel-Driven-Defense.pdf

> Once you understand where most of your team's time and energy is being spent you can start to close technology gaps and automate mundane and manually intensive parts of the incident handling process.

features, integration points and automation capabilities that are not already being used. If they fall short, a purchase can often be justified using existing metrics that quantify risk and operational costs.

Automate workflows

To effectively detect and respond to attacks, security analysts must gather and analyse multiple complementary data sources as one. This includes network traffic, endpoint data (including mobile devices), application data and logs. They also need real-time threat intelligence from internal and external sources to detect and validate an incident so they can respond appropriately. Automation is key to reducing the metrics that matter. The following examples illustrate the benefits of workflow automation.

Example 1: Threat Intelligence for Alert Validation

Before

To validate alerts, analysts manually pivot on datasets in the SIEM. They collect data from network forensics products and physically access the affected machine to run live response scripts. The average time to validate an alert is **4 hours.**

After

Information about the threat gathered from next-generation threat-detection products can be converted into machine-readable threat intelligence and used to automatically check the suspected endpoint for alert validation. The average response time for these alerts is reduced to **15 minutes.** When there isn't enough information to automatically validate an alert, live response data, including recorded process activity, is automatically collected. Preliminary analysis steps are automatically conducted, such as malware analysis and populating data points with contextual information from threat intelligence sources, to provide analysts with the necessary data to validate the alert and begin answering important questions. The average time to validate is **30 minutes**.

Example 2: Live Response Data for Forensic Analysis

Before

Analysts require a significant amount of time to acquire data and perform analysis to reconstruct the incident and understand how the attacker gained entry and what they did. Network forensics experts spend **4 hours** analysing sessions to understand the initial compromise, command and control traffic and what was transmitted. Meanwhile, forensic images of disk and memory are made on-site, taking **8 hours** or more to complete per affected system. Forensics experts — relying solely on dissecting the forensic images — take an additional **8 hours** to reconstruct attacker activity on each system. The average time to complete analysis of each system is **16 hours** from the time the incident was verified.

After

Live response data collection, including recorded process activity, is already available and the preliminary analysis is complete. The security analyst has a good picture of what happened. He or she can review the recorded process activity before pivoting to review disk and memory forensic data to fill in missing details. The analyst then retrieves additional data remotely as needed such as interesting files on disk. At the same time they are able to view related network session information and peer into content, including commands sent by the attacker and transmitted data. The average time to analyse each affected system has been reduced to **2 hours** from the time the incident was verified.

Step 3: Detect

Advanced, targeted attacks are not instantaneous events. They involve a series of actions and multiple phases that occur over a period of time. Professional cybercriminals are so adept at cloaking their activities that they routinely go unnoticed for months and often years. They conduct detailed reconnaissance activities and when necessary, develop custom-tailored exploits to penetrate your enterprise network and steal sensitive corporate data, intellectual property, business plans and personal information.

Detecting security incidents early in the attack lifecycle is therefore paramount to success and lowers the complexity and associated cost compared to detecting after the organisation has been compromised.

Figure 1: The Attack Lifecycle

Initial Compromise	Establish Foothold	Escalate Privileges	Lateral Propagation	Data Staging & Exfiltration
Gain Initial Access	**Strengthen Position**	**Steal Valid User Credentials**	**Access Other Servers & Files**	**Package & Steal Target Data**
Phishing e-mail Watering-hole attack Removable media Malicious download	Custom malware Command and control 3rd party application exploitation	Credential theft "Pass-the-hash"	Window & Linux lateral movement techniques Reverse shell access	Staging servers & directories Data consolidation Data theft

Threat detection methods

Effective threat detection requires security teams to employ multiple techniques simultaneously. This includes monitoring Internet traffic, endpoints and common initial infiltration vectors. Below are overviews and guidance on popular threat-detection methods that can help security teams identify an attack during the initial phases of the attack lifecycle.

Traditional tools

Despite claims to the contrary, traditional tools such as antivirus, network IDS, firewalls and web application firewalls are not relics of the past. Though not bulletproof, these tools act as a front line of defence to reduce noise and can trigger alerts even during sophisticated cyberattacks. Because these alerts don't always provide enough context to guide a proper response, we recommend augmenting these alerts with data from other sources.

Machine-readable threat intelligence

Machine-readable threat intelligence is a cornerstone of detection. It takes signatures to a new level. Sources of threat intelligence include:

- Shared indicators of compromise (IOCs) released from vendors and government agencies that are codified by threat intelligence standards (such as STIX, OpenIOC) and contain:
  - MD5 hashes
  - Filenames and sizes
  - DNS hostnames
  - IP addresses
  - File and registry modifications
  - Malware attributes
  - Running process mutexes
  - PE header information
  - SNORT rules
  - YARA rules, and more
- Commercial and open source threat intelligence feeds
- Indicators shared by industry associations and industry-specific information-sharing groups (such as ISACs, InfraGard and ISSA) that privately share threat intelligence
- Internal threat intelligence from your own incidents that you curate, document and maintain in your own threat intelligence library

Applying threat intelligence

Having machine-readable threat intelligence is not enough. It must be made actionable by applying it to available data sources or points of visibility (e.g. endpoints, network traffic, log files and application data). Solutions are available that monitor process activity in real-time or poll the system state and record process activity at intervals that can span weeks. The closer that monitoring is to real time, the better. Ensure that selected systems can ingest multiple machine-readable threat intelligence formats and can integrate with threat feeds either directly or via a normalisation engine like the Collective Intelligence Framework (CIF).

When monitoring endpoints, there's a balancing act between low detection times and the impact on endpoint performance. Being able to throttle that performance impact is crucial to minimise business disruption. In the event of a critical security incident, adjust throttling as priorities shift.

For network traffic, consider the performance impact of placing sensors in-band versus out-of-band. Monitor for sensor CPU spikes and packets being dropped during peak hours. Although more hardware helps, there can be a drastic difference in resource utilisation from one vendor to the next. For centralised log files, a low time to process input and generate alerts is vital.

Behavioural

Also known as *heuristics, signatureless* and *next-generation threat detection*, behavioural profiling analyses network communications to identify abnormal activity that may indicate that an attacker is present on your network. In the case of malware, a score is assigned to various features and actions of executables and running processes. Malware-specific behaviours receive a high score while features that are also present in innocent executables will receive a lower score. The higher the cumulative score on a binary, the more likely it is to be malicious.

Attackers often repeat behavioural patterns, such as:

- Naming conventions
- Working directories used to copy files
- Methods of using built-in system commands and utilities

By looking for these patterns, analysts can identify an attack in progress even if no malware is present. Additionally, attackers usually use compromised credentials. Applying heuristics to user behaviour, and especially to administrative and service accounts, helps to identify suspicious behaviour based on the systems they authenticate to, network traffic and application access.

Security analytics

Analytics is the practice of pivoting and filtering large, aggregate datasets to identify anomalous items or events that raise suspicion. Frequency analysis is one popular method. On endpoints, analysts can collect a list of running processes and associated metadata, such as what they look like on disk and open network connections. Oddities will stand out when sorted on columns such as file path. The same method can be employed against installed drivers and services, persistence mechanisms/autoruns, etc. By researching entries with a low number of occurrences, new malware or variants missed by threat-detection products and threat intelligence can be discovered. We recommend that you also:

- Investigate spikes in network traffic occurring at unusual times or from sources that don't produce much network traffic.
- Monitor hosts with defined functions that produce reliable traffic patterns for deviations. Within the network traffic, there may be unique attributes as it's parsed, like an unusual user agent string used by malware.

By identifying and investigating these anomalies, investigators can more easily identify the proverbial needle in the haystack. Further, products can combine analytics with heuristics to help automate threat detection. Outliers with a high threat score are worthy of investigating.

Service providers

There are a growing number of service providers that will monitor your network for threats. These service providers will often use proprietary threat intelligence to monitor your network traffic, log files and endpoints. These services can act as an excellent way to improve threat detection, especially if in-house capabilities are in process of being built out. When forming relationships with service providers, pay close attention to where they have points of visibility and their service level agreements. These will impact both their ability to detect threats and your key detection and response metrics. Review the output to determine if they send a bare bones alert or if it includes preliminary analysis results. With most service providers, the alert and possibly an initial triage will be the extent of the work provided. Additional analysis and incident response services are usually charged by the hour. *Appendix D, Product Requirements*, can be used to evaluate service providers to understand how they can help fill gaps to achieve rapid detection and response.

Step 4: Respond

During the Respond step, security teams confirm, analyse and document attacks that they have detected in the previous phase. The goal is to assess the impact so they can develop the appropriate strategy to remediate and resolve the incident. This is where most organisations face severe challenges, including poor metrics for response and remediation. What follows is an overview of the typical response steps along with guidance on how to perform them quickly. For in-depth reading on incident response, we recommend "Incident Response & Computer Forensics," Third Edition (ISBN-10: 0071798684).[4]

Validate

To triage an alert and determine whether it is a false positive or a valid threat, analysts can:

- Use leads in the alert (IP addresses, DNS hostnames, machine names and timestamps)
- Pivot to view related SIEM information in the SIEM
- Review netflow data and live response data from the suspected endpoint

To reduce noisy false positives and redundant alerts, a best practice is to modify the responsible systems, rules and indicators. Often, a simple modification of an indicator is enough to correct the problem. For more information, refer to the blog post by the author of "Applied Network Security Monitoring" (ISBN-10: 0124172083),[5] "Calculating IDS Signature Precision."[6]

If an alert includes enough useful details, you can codify it into machine-readable threat intelligence and match it against endpoint data and other relevant data sources to automate the validation process. For other alerts, automate the collection of data needed by analysts. You can automatically enrich the data you collect with contextual information from threat intelligence sources by looking up IPs, hostnames and file hashes. This enables you to highlight known threats in the results.

Frequently, malware detection products constitute the largest percentage of alerts. During validation, it will become obvious which family the malware belongs to and how that family of malware is used. If the malware family presents a relatively low-risk, like adware or botnets, it may not be worth the time to perform a comprehensive analysis unless there is additional information that indicates there is a serious issue. For these nuisance threats, consider foregoing comprehensive analysis in favor of containment and generic remediation as a trade-off for the minor risk.

4 Luttgens J., "Incident Response & Computer Forensics, Third Edition," 2014 [online] http://www.amazon.com/Incident-Response-Computer-Forensics-Edition/dp/0071798684

5 Sanders C., "Applied Network Security Monitoring," 2014 [online] http://www.amazon.com/Applied-Network-Security-Monitoring-Collection/dp/0124172083/

6 Sanders C., "Calculating IDS Signature Precision," 2014 [online] http://www.appliednsm.com/calculating-ids-signature-precision/

Contain

If the validated threat warrants further investigation, the affected endpoint should be contained. Depending on your organisation's risk tolerance, this could mean containing network communications so that the endpoint can only communicate with incident response platforms. Another option is to take a hybrid approach by cutting off Internet connectivity to allow the endpoint to continue functioning on the intranet. These steps can be automatically performed by locking down communications at the endpoint or network equipment once you validate a threat. For rapid detection and response, automatic containment triggered by validation is recommended. However, an analyst can also trigger containment by manually initiating the process.

For some organisations, there may not be executive support to contain endpoints at this stage. By keeping track of other metrics and documenting lessons learned in the Resolve step, you can calculate the risk and cost associated with delaying containment. Maintaining accurate asset management and inventory, including system owner, role, location, etc., can go a long way in building executive support by explaining very specific rules that will trigger automatic containment.

Best practice tip:
Record the time between validation and containment and track this time across incidents as a time-to-contain metric.

Collect

If a validated threat warrants further investigation, you will need to collect data to answer the questions presented below. *Appendix A, Threat Activity*, provides a worksheet that you can use to record this information.

1. What is the nature of the threat?
2. How and when did the threat enter the system it was detected on?
3. What actions were performed by the threat since it entered the system?
4. Were any credentials compromised?
5. Were additional backdoors planted or created?
6. Did the threat move laterally to other systems?
7. Are there signs of data theft such as copies of documents on the system or unusual spikes in outbound network traffic associated with attacker activity?

Data sources include netflow data, captured network sessions, log files, endpoint forensic data, recorded process activity and application data. Anything that cannot be quickly queried on demand will need to be collected. By automatically collecting the data needed for analysis you can significantly improve your incident response metrics.

For endpoints, collecting live response data is common practice. Live response data consists of useful endpoint metadata that can be quickly obtained to triage an alert without going through the lengthy process of creating and managing full disk and memory images. Traditionally, live response collections contain system state information such as installed software, running processes, open network connections, file and registry listings and programmes configured to survive a reboot.

Best practice tip:
Record the time between containment and collection and track it across incidents as a time-to-collect metric. If containment wasn't performed, measure the time between detection and collection.

More recently, endpoint detection and response (EDR) tools have made it increasingly popular to query endpoint process activity, which the EDR tools record. Specific information that EDR tools collect includes process execution, network activity, file modification and registry modification. Analysts can query and review recorded activity to reconstruct what happened without relying solely on traditional forensic data, which is more difficult to piece together and often incomplete. Be sure to take the deployment architecture into consideration when reviewing solutions that record process activity. Centrally logging and analysing process activity from tens or hundreds of thousands of endpoints presents its own challenges. Securely storing the activity logs on the endpoints themselves is a more efficient approach with the tradeoff in risk that an attacker could theoretically delete or alter recorded data.

Analyse

Once the data has been collected, the security analyst or incident responder can reconstruct what happened to answer the important questions listed previously. Various techniques are available.

Timelining

Alerts usually contain useful timestamps. A timeline combines multiple data sources such as file and registry timestamps, event logs, recorded process activity and network traffic into a single view, sorted chronologically. By reviewing entries around the time of interest, you can identify and trace other activity related to the incident. As you discover new indicators attributable to the threat, filter the display and pivot to adjust what is seen.

Autoruns

A backdoor can survive a reboot through the use of registry keys, installed services and drivers, browser helper objects (BHO), etc. Reviewing a list of autorun entries with associated metadata is an excellent way to identify malware and backdoors planted by attackers.

Binary analysis

Understanding the capabilities and properties of executable files and running processes is necessary when seeking unknown malware. The two most common techniques are static analysis and dynamic analysis. Static analysis looks at the file as it exists on disk. Dynamic analysis runs the file and records activity. Further analysis can be performed by automatically or manually disassembling the binary to look at the code forks and how they operate. All three analysis methods are valuable. Due to the high cost and time for reverse engineers, it is best to use an automated system that can provide useful results. Using sophisticated heuristics to apply threat scoring helps to focus in on what you're looking for more quickly.

Network forensics

As network-based indicators are discovered, such as malicious IPs and DNS hostnames, analysts can perform network forensics on captured sessions to reconstruct the commands and data transmitted. As SSL/TLS and SSH have risen in popularity, analysts should ideally have the ability to decrypt and review that traffic to reduce blind spots. The ability to view content, such as web pages and email, and extract interesting binaries will speed up the investigation considerably.

As you perform your analysis you will likely need to acquire more data, including specific binaries from disk or even full images. It is important that the analyst be able to do this quickly from within the incident response tool. Because incidents may lead to litigation, it is important to acquire and store the data in a forensically sound manner and adhere to forensics practices, such as tagging assets and documenting chain of custody.

How to speed up the review process

Recognising that there are volumes of data to sift through, your solution should:

- Have lists of known good and known bad to filter with. This should include the ability to create custom lists as environments are baselined and incidents are investigated.
- Integrate threat intelligence sources so it automatically applies machine-readable threat intelligence, highlights matches and provides context. The ability to automatically provide information from threat intelligence sources by looking up IPs, hostnames and file hashes will also drastically speed up analysis time.
- Integrate with other analysis engines in use so that samples can be submitted and the results incorporated with the data.
- Enable pivoting between endpoint data and network data. Because these two data sources are related, much time can be saved by enabling analysts to easily follow threat activity through both sources simultaneously.

Step 5: Resolve

Incident resolution involves removing threats, restoring services and applying lessons learned from the incident to bolster preventive defences and improve ongoing rapid detection and response. To eradicate the threat, you must develop and execute a plan to remove planted backdoors, fix exploited vulnerabilities and reset compromised user credentials. All actions should be logged to facilitate recovery reporting. Automating common remediation actions and applying them across both the network and endpoints significantly accelerates these processes and saves valuable human time.

In the event that an endpoint was compromised by an attacker moving laterally from another system, you should take care to quickly identify and resolve all compromised endpoints. When faced with sophisticated threat actors, there is a risk that they may notice before resolution is complete and take steps to thwart the effort. We therefore frequently advocate that organisations scope all points of entry across the enterprise and close them all at once to avoid a game of "whack-a-mole." The decision isn't an easy one to make and involves understanding the risks presented by both approaches. Each organisation must decide which approach to use based on their risk tolerance and the specific circumstance.

Best practice tip:
Record the time between analysis and resolution and track it across incidents as a time-to-resolve metric.

Incidents deemed to be mass malware, like spyware and botnets, can be resolved by manually cleaning the system or reimaging. Submitting these undetected malware samples to antivirus vendors will help improve preventative defences and reduce recurrences. However, you should retain malware used in targeted intrusions to ensure your intelligence isn't made known to the attacker. Document threat intelligence gained during your analysis, store indicators of compromise as machine-readable threat intelligence and add it to the threat library to improve your detection and analysis capabilities.

Close the RDRM loop

In order to improve your rapid detection and response capabilities, it is important to circle back to the Identify step to close the feedback loop. This will enable you to apply lessons learned from incident handling to prioritise preventative defences and identify weaknesses hindering rapid detection and response. Actions to take include:

- Create new technical requirements for detecting and responding faster.
- Automate manual steps.
- Prioritise preventative defences based on initial threat vector and weaknesses that allowed the threat to make progress. For example, if attackers were able to steal and crack weak Windows LM hashes to escalate privileges, disable the use of LM hashes to make privilege escalation more difficult. Or if attackers repeatedly access their first victim by exploiting an outdated application, give patch management more attention.
- Adjust what data is acquired for analysis during the Validate and Collect activities in the Response step by paying attention to cost/benefit. Is the time it takes to gather specific data worth it?
- If the time-to-detect was lengthy, identify weaknesses in the detection capabilities. For example, it is possible that analytics aren't being performed frequently enough or there aren't enough quality threat intelligence sources.

Conclusion

In just a few years, the motives and tactics of attackers have rapidly evolved to include economic and industrial espionage, cyberwarfare, organised crime, hacktivism and terrorism. It's become widely accepted that security incidents are inevitable for any organisation that has valuable data. To defend against determined attackers, organisations must accelerate their ability to detect, investigate and stop attacks.

Rapid detection and response is not a new concept. It's been done by leading SOCs and IR teams for years through tremendous in-house efforts with dedicated programmers to integrate and automate a multitude of disparate point products. Thankfully, the security vendor ecosystem has been moving in the direction of consolidating and integrating complementary capabilities, making rapid detection and response technologies more accessible.

As organisations struggle to overcome talent shortages, keep up with modern threats and reduce risk, efficiency has become a necessity. The stakes are too high and there simply aren't enough skilled people to continue relying on overworked, scarce experts. We believe every organisation is capable of using the RDRM to disrupt attack lifecycles and achieve a faster and more effective incident response that comes from greater visibility and context, consolidation and integration of security tools and automation of mundane steps.

How Fidelis Cybersecurity Can Help You

Every day security teams face well-funded adversaries. Inspired by a range of motives, these attackers have a single mission — to steal intellectual property, customer information and sabotage critical infrastructure. Fidelis products and services prevent attackers from achieving this mission. By focusing on real-time detection, prevention and continuous response, Fidelis empowers organisations to reduce the theft of assets and data, improve ROI on security investments, improve the efficiency of security analysts, lower incident response costs and reduce reputation risk.

Proactive, advanced threat defence

Attackers hide their exploits deep inside your network, email and endpoints. Fidelis digs as deep as attackers hide. Our products enable organisations to reduce the time it takes to detect and resolve incidents, prevent data theft and stop attackers at every stage in the attack lifecycle.

- **Visibility Across the Attack Lifecycle.** We don't just look for the tools attackers use. We also look at their behaviour. When we find an attack, we provide visibility that enables you to reconstruct attackers' footprints so you can see where they have been — even when they are accessing devices, such as laptops or mobile devices, which have left your network.
- **Network + Endpoint + Mobile.** We leave attackers no place to hide. Our products detect attackers on the network across all ports and protocols. Then, we pursue them out to the endpoints where your data lives.
- **Pivot from Detection to Investigation.** We are a rare breed. We don't make work for you with a deluge of alerts. We save you time by allowing you to move from alert to investigation with a single click and in a single solution.

Fidelis consulting services — the power of experience

Fidelis' security consultants have decades of experience assisting organisations of all sizes to prepare for and respond to security incidents. The Fidelis Security Consulting team has the scale, experience and credibility you need. In addition to unparalleled expertise in defence against cyber warfare, theft of intellectual property, advanced persistent threats and other malicious threat actors, our consultants bring a commitment to confidentiality to protect reputations and valuable brands.

We offer a full portfolio of consulting services designed to enable you to effectively and efficiently respond to and completely recover from a cyber attack:

- **Incident response services.** Immediate assistance to determine the scope of the incident, remove attackers from your environment and re-secure your network.
- **Proactive security assessments.** Understand how your security programme compares to your peers and industry best practices. Based on your needs our experts will evaluate, assess and validate your incident response plan or your entire security programme.
- **PCI services.** Fidelis is certified as a Qualified Security Assessor by the PCI Security Standards Council to validate the adherence to PCI Data Security Standard (DSS).
- **Security Operation Centre (SOC) development services.** Fidelis can build your organisation's security operations centre (SOC) from the ground up or help mature your existing SOC.
- **Litigation support services.** Fidelis security professionals possess forensic certifications and are experienced in collecting, analysing and preserving data required for depositions, investigations, discovery and testimony.

Appendix A. Threat Activity

THREAT ACTION	ANSWER
What is the nature of the threat? Threat classification is useful for understanding severity and identifying trends.	
How and when did the threat enter the system it was detected on? Cataloging the initial infiltration vector helps identify and prioritise exploitation paths and can uncover larger breaches in cases where a system is compromised through lateral movement from another compromised system.	
What actions were performed by the threat since it entered the system? Reconstructing attacker activity is the only way to uncover compromised credentials and stolen data; and identify all points of entry available to the attacker.	
Were any credentials compromised? Even if backdoors are removed, stolen credentials can provide attackers with access to VPN, email and other points of entry exposed to the Internet.	
Were additional backdoors planted or created? Attackers commonly drop backdoors of varying families, both passive (e.g. webshells) and active (e.g. RATs) on several systems to ensure uninterrupted access to the company intranet.	
Did the attacker move laterally to other systems? Attackers commonly gain entry to one system and then pivot using stolen credentials to connect to other systems.	
Are there signs of data theft such as copies of documents on the system or unusual spikes in outbound network traffic associated with attacker activity? Attackers often copy stolen data to other systems used as a staging area, then upload the data, blending in with normal web traffic outside the network.	

Appendix B. Rapid Detection and Response Metrics Self-Assessment

HOW LONG DOES IT TYPICALLY TAKE TO	ANSWER
Discover threats that bypass defences?	
Validate security alerts?	
Contain affected endpoints?	
Collect the data needed for analysis (endpoint, network, logs and application)?	
Conduct analysis and identify the threat, root cause and scope of the incident?	
Resolve the incident?	
What percentage of incidents would you guess go undetected until obvious symptoms appear?	
What percentage of incidents is thoroughly investigated to answer important questions?	
Do these numbers change for remote workers and satellite offices?	

Appendix C. Rapid Detection and Response Model

PROCESSES			
For each incident detection mechanism, there are established procedures that address the following areas:	**Current Capability (0-10)**	**Desired Capability (0-10)**	**Priority Score (0-10)**
Metrics are maintained for security alert and incident handling times. Metrics are kept for incident time-to-detect (dwell time), time-to-validate alerts, time-to-contain, time-to-collect, time-to-analyse and time-to-resolve. These times help identify inefficiencies and risk. Response management and ticketing systems should be used to automatically collect and provide this information in a standard time zone.			
Weaknesses exploited for initial entry and attack progression are documented. By understanding how threats enter and progress, organisations can prioritise and drive security initiatives.			
False positives are documented and used to tune threat detection systems. Reducing noise saves valuable time and helps prevent analysts from ignoring alerts from security products.			
Tabletop exercises, such as the one made freely available by FEMA, are conducted on a regular basis.[7] Tabletop exercises will expose additional requirements for technology and processes currently lacking and prepare organisations for a serious breach so mistakes are less likely to happen.			
ENDPOINT VISIBILITY			
For each incident detection mechanism, there are established procedures that address the following areas:	**Current Capability (0-10)**	**Desired Capability (0-10)**	**Priority Score (0-10)**
Low-level forensically sound access to disk drives, removable media and computer memory An agent with kernel-mode access is required for deep forensic visibility. While there are ways to obtain data without an agent, severe limitations hinder usefulness and the data returned is commonly manipulated by running malware.			
Parse the file system of disk drives and removable media Being able to read the disk raw and parse the file system is necessary to identify hidden malware and gain better context through metadata not accessible using operating system APIs.			

7 FEMA. National Level Exercise 2012: *Cyber Capabilities Tabletop Exercise*, 2012 [online] http://www.fema.gov/media-library/assets/documents/26845?fromSearch=fromsearch&id=5949

ENDPOINT VISIBILITY *(continued)*			
For each incident detection mechanism, there are established procedures that address the following areas:	**Current Capability (0-10)**	**Desired Capability (0-10)**	**Priority Score (0-10)**
Parse the Windows registry Being able to read and parse registry files is necessary to identify hidden malware persistence entries.			
Parse operating system logs Operating system logs contain information with timestamps useful for reconstructing an incident and identifying lateral movement throughout a network.			
Parse web browser history Web browser history can play an important role with insider threats as well as identifying the root cause of malware infections. Make sure multiple common web browsers are supported.			
Parse memory as used by the host operating system Malware commonly obfuscates and hides itself on disk and network communications. Identifying suspicious processes, understanding their capabilities and authoring indicators of compromise becomes easier with memory forensics.			
Obtain metadata from disk, removable media and memory Being able to rapidly collect metadata for triaging is essential when triaging alerts and confirming incidents.			
Obtain specific binaries (files or otherwise) from disk When reviewing metadata, there is often the need to acquire specific files from disk, such as unknown executables and files created by attackers.			
Obtain forensically sound images of disk drives and memory Sometimes it is necessary to obtain full images in a standard container format to use with multiple tools in order to uncover crucial attacker activity, including the files they've created and modified. Forensic images will also be required in incidents that will go through litigation.			
A recorded history of process execution, network activity, files accessed and modifications made to the system Logging process activity at each endpoint allows security analysts to validate security alerts and perform an initial analysis more quickly than piecing together forensic data that doesn't tell the whole story.			
Record inserted removable media information and files copied Being able to identify and investigate threats presented by removable media, such as malware and data theft, requires visibility into these devices.			

ENDPOINT VISIBILITY *(continued)*			
For each incident detection mechanism, there are established procedures that address the following areas:	**Current Capability (0-10)**	**Desired Capability (0-10)**	**Priority Score (0-10)**
System containment support When a system is confirmed as compromised, it is advisable to limit its network communications to prevent an attack from progressing while data collection and analysis takes place.			
Remediation actions The ability to kill processes and delete files speeds up remediation tasks.			
Upload and execute For complex sequences of actions that need to be performed or to run additional tools, it is advisable to have the ability to upload and execute applications, installers and scripts.			
Data leakage detection Identifying sensitive data where it shouldn't be is important for both detecting risky user behaviours and for identifying the latter stages of the attack lifecycle as an attacker attempts to stage and exfiltrate stolen data.			
Agent just-in-time deployment For systems where agents are not permitted to run continuously, the ability to push an agent on demand manually or through automation is needed when searching for indicators of compromise or collecting data for analysis. Once the task is performed, the agent should be capable of self-dissolving.			
Remote agent support Communicating with systems off the corporate network or at remote sites, including those behind NAT, is needed for true enterprise-wide visibility.			
Throttling support Limiting CPU and network resource consumption provides control over system performance degradation and the saturation of network WAN links.			
Job resumption support For lengthier tasks, such as creating images, collecting several binaries or performing deep analysis that takes time at the endpoint, it is important that interrupted jobs are able to resume as systems go offline and come back online.			
LDAP support for targeting computer and user groups On Windows networks, the ability to specify systems and users in Active Directory is valuable to limit where tasks are performed. Additionally, it allows policies like process activity recording to be configured based on organisational structure.			

ENDPOINT VISIBILITY *(continued)*			
For each incident detection mechanism, there are established procedures that address the following areas:	**Current Capability (0-10)**	**Desired Capability (0-10)**	**Priority Score (0-10)**
Support for desktop and server OSs in use Although there are more attackers familiar with Windows networks, there are plenty that know how to exploit Linux, Solaris, Mac OS X and embedded OSes such as those used in point-of-sale environments.			
Mobile OS support Visibility into what is happening on mobile these devices is needed since they have user credentials and access to sensitive data.			
Monitor endpoint activity and conduct scans using machine-readable threat intelligence Threat intelligence must be applied at the endpoint level to identify activity and forensic data related to attacker activity that isn't seen from the network perspective or captured in centralised logs. The closer to real-time the monitoring can get the better although some threat intelligence may require periodic scans to avoid system performance degradation.			
Behavioural malware detection The characteristics of executable files and their running behaviours can be analysed to identify traits commonly attributable to malware. A risk score can be calculated to identify samples bypassing preventative defences undetected.			
Security analytics support Review enterprise-wide datasets with flexible filters and pivot for the purpose of conducting analytics.			
A flexible data analysis interface Being able to pivot, filter, sort, search and tag data under review allows analysts to cut through the noise and follow leads.			
Query all systems for specific data sets or to perform checks and receive results in a timely fashion Enterprise visibility requires the ability to query all systems in a timely fashion to identify indicators of compromise, perform analytics and scope a compromise when multiple systems known and unknown are affected.			

NETWORK VISIBILITY			
For each incident detection mechanism, there are established procedures that address the following areas:	**Current Capability (0-10)**	**Desired Capability (0-10)**	**Priority Score (0-10)**
Sensors placement Sensors should be placed at Internet egress points at a minimum, optionally extended to inter-segment traffic for critical network segments.			
Deep session inspection Being able to decode and analyse content in real-time, no matter how deeply embedded, is vital to gain visibility into applications and in particular, the content that is flowing over the network.			
Identify attacker behaviour To unilaterally block unauthorised transfers of information in real-time, across all ports and protocols, without depending on third-party proxies, analysts must be able to identify attackers when they are active on the network.			
Content-rich metadata support In order to review network activity, data more verbose than basic netflow information is needed. Detailed metadata that provides a rich historical view of all network communication — protocols, applications and content — is needed to provide the context required to understand events taking place in the network.			
Network session reconstruction The ability to follow specific network sessions is invaluable when honing in on sessions of interest.			
Content rendering (e.g. web pages, emails, images) Being able to quickly see the rendering of a web page or email inside a packet capture allows analysts to quickly review sessions for threat activity.			
Filtering Being able to sift through larger data sets based on specific datapoints, such as IPs and DNS hostnames, is needed to scope out affected traffic and systems.			
Binary identification, extraction and analysis In a packet capture, there will be files copied and malware delivered. Pulling files out of the network stream for further analysis is a common analysis step.			
SSL/TLS and SSH decryption It is important to be able to decrypt encrypted sessions to detect and analyse threats from the network perspective. An increasing amount of web traffic is delivered over SSL/TLS and attackers are encrypting their command and control infrastructure.			

NETWORK VISIBILITY *(continued)*			
For each incident detection mechanism, there are established procedures that address the following areas:	**Current Capability (0-10)**	**Desired Capability (0-10)**	**Priority Score (0-10)**
Monitor traffic against internally maintained and externally supplied machine-readable threat intelligence Being able to make machine-readable threat intelligence actionable by monitoring traffic for network-based indicators such as IPs, DNS hostnames and URLs is a must-have.			
Analytics capabilities with flexible filters, pivots and visualisations By reviewing larger network traffic data sets to identify anomalous activity and interpret that into visualisations, unknown threats can be identified.			
LOG MANAGEMENT			
For each incident detection mechanism, there are established procedures that address the following areas:	**Current Capability (0-10)**	**Desired Capability (0-10)**	**Priority Score (0-10)**
Centralised logging Critical system, network device and application logs are centralised and preserved for a minimum of one year. This includes VPN, web applications, web-based email, DHCP, DNS, firewall and proxy.			
Endpoint logs Endpoint logs that cannot be collected due to resource constraints (e.g., end-user workstations) should be configured to log the same events with a high retention period.			
Logging configuration Logs are configured to capture events critical to reconstructing attacker activity such as netflow, authentication success and failure and process execution.			
THREAT INTELLIGENCE			
For each incident detection mechanism, there are established procedures that address the following areas:	**Current Capability (0-10)**	**Desired Capability (0-10)**	**Priority Score (0-10)**
Threat actor profiles Maintain profiles on known threat actors with deep intelligence on how they operate, including their tactics, techniques and procedures (TTPs).			
Service providers Service providers can monitor log files, network traffic and endpoints to identify threats using proprietary threat intelligence.			

THREAT INTELLIGENCE *(continued)*			
For each incident detection mechanism, there are established procedures that address the following areas:	**Current Capability (0-10)**	**Desired Capability (0-10)**	**Priority Score (0-10)**
Machine-readable threat intelligence formats Multiple formats should be read and made actionable at endpoints, network traffic and centralised logfiles.			
OpenIOC A flexible format for documenting indicators of compromise, created by Mandiant and used by some information sharing groups.			
STIX A standard created by MITRE for the U.S. to document threat intelligence, currently being adopted by ISACs.			
YARA YARA rules are designed to identify malware and used by many existing products.			
Delimited files (e.g. csv) Many threat feeds and sharing groups still use plain, delimited files.			
Machine-readable threat intelligence sources			
Internal When investigating security incidents, any observable attacker activity and tools analysed should be codified and put to use to detect other occurrences and future attacks.			
Commercial Several vendors provide threat intelligence feeds and lookup services to add additional context. Multiple vendors should be used depending on their strengths and relevance to your organisation.			
Information sharing Information sharing groups such as ISACs and InfraGard are highly valuable since the threat intelligence tends to focus on relevant threat actors targeting organisations.			
Open source Blog posts, whitepapers, FBI bulletins and other sources are freely available on the Internet. Sometimes they provide machine-readable threat intelligence. Other times, you will need to manually create it.			

Appendix D. Product Requirements

For each incident detection mechanism, there are established procedures that address the following areas:	Current Capability (0-10)	Desired Capability (0-10)	Priority Score (0-10)
Role-based access control (RBAC) Access to security systems should provide user and group management with the ability to limit access to specific groups to achieve least-privileged access based on the role of the user.			
Secure communications Communications for management, inter-component interaction and API access should all be encrypted and secured against interception and modification.			
Recordable metrics By keeping track of incident response metrics in security systems, it becomes much easier to spot trends and gaps in rapid detection and response capabilities.			
Alerting methods Alerts should be available in a dashboard view, through standard alerting protocols such as CEF, and through exported reports.			
Existing integrations There should be integrations into complementary products to provide greater context and automate workflows.			
SIEMs Automatically contain affected systems, collect data to display in the SIEM, perform preliminary analysis and provide actions that can be triggered by security analysts from the SIEM interface.			
Next-gen threat detection appliances Automatically validate alerts, contain the threat, collect data and perform preliminary analysis.			
Sandboxes Submit suspicious files to sandboxes and incorporate the results into the analysis interface.			
Programmability An API should be available for both input and output to enable integration and extend automation.			

Fidelis Cybersecurity is creating a world where attackers have no place left to hide. We reduce the time it takes to detect attacks and resolve security incidents. Our Fidelis Network™ and Fidelis Endpoint™ products look deep inside your traffic and content where attackers hide their exploits. Then, we pursue them out to your endpoints where your critical data lives. With Fidelis you'll know when you're being attacked, you can retrace attackers' footprints and prevent data theft at every stage of the attack lifecycle. To learn more about Fidelis Cybersecurity products and incident response services, please visit **www.fidelissecurity.com** and follow us on Twitter **@FidelisCyber**.

 WP_Fidelis_RDRM_UK_1603