LMNTRIX

BE THE HUNTER | NOT THE PREY

HOW **MACHINE LEARNING** & UNDERGROUND INTELLIGENCE AUGMENT **THREAT INTELLIGENCE**

Staying ahead of advanced persistent threats using proactive defense strategies

LMNTRIX USA.
333 City Blvd West,
18th Floor, Suite 1805
Orange, CA 92868
+1.888.958.4555

LMNTRIX UK.
200 Brook Drive, Green Park,
Reading, RG2 6UB
+44.808.164.9442

LMNTRIX INGAPORE.
60 KAKI BUKIT PLACE#05-19
EUNOS TECHPARK
+65 3159 0639

LMNTRIX HONG KONG.
14F, Manning House, 38-48
Queen's Road Central, Central,
Hong Kong
+852.580.885.33

LMNTRIX AUSTRALIA.
Level 32, 101 Miller Street,
North Sydney NSW 2060
+61.288.805.198

LMNTRIX INDIA.
VR Bengaluru, Level 5,
ITPL Main Rd,
Devasandra Industrial Estate,
Bengaluru, Karnataka 560048,
Email: sales@lmntrix.com
+91-22-49712788

WHITE PAPER **2023**

lmntrix.com

LMNTRIX

BE THE HUNTER | NOT THE PREY

EXECUTIVE **SUMMARY**

Cyber security is a constant battle of wits between organizations and attackers as they look to gain an advantage. Organizations implement security controls to counter the latest attack methods, but these methods evolve to outmaneuver these controls. Attackers then develop new threats, seeking exploitable vulnerabilities and creating novel malware.

Gaining a thorough understanding of the threat landscape that an organization inhabits allows the accurate identification, prioritization, and mitigation of risks. Consequently, knowing what credible threats are likely to target the organization means they can apply the best protective measures to counter them.

Using threat intelligence can provide organizations with valuable information about potential security threats and help them protect against cyberattacks. Technical and tactical threat intelligence can enable security teams to counter current threats and uncover previously unknown attacks. Operational threat intelligence enables assessing immediate threats to drive short-term planning and decision-making. Finally, strategic threat intelligence can allow long-term planning, improve security posture, and inform security policies and procedures to maintain security in an evolving threat landscape.

Threat intelligence offers a proactive means to counter the latest threats and predict future threats to gain the upper hand in the battle against attackers. This technique analyzes data related to security events, incidents, and the attackers' actions. Understanding the motivation, targets, and behavior of attackers helps detection of new threats, respond to existing attacks, and forecast future protective measure requirements. A key source of threat intelligence is information only available from the dark web, the world wide web area outside the reach of standard search engines. Here, criminals discuss threats and vulnerabilities and share intelligence on targets, trade tools, and information on illicit marketplaces.

The critical challenge for threat intelligence processes is the handling and analyzing vast quantities of data of varying quality, relevance, and usefulness to extract the needle in the haystack that is actionable information. Machine learning techniques offer a practical, scalable solution for managing these processes accurately, quickly, and efficiently.

Augmented threat intelligence then combines the artificial intelligence-based functions with the existing security solutions to deliver a complete integrated solution that actively supports the organization's security team to enable them to work better.

This paper looks at how these augmented threat intelligence program elements operate and how LMNTRIX leverages this solution to deliver enhanced security protection to our clients.

LMNTRIX

BE THE HUNTER | NOT THE PREY

CONTENTS

LMNTRIX

BE THE HUNTER | NOT THE PREY

LMNTRIX

BE THE HUNTER | NOT THE PREY

THREAT INTELLIGENCE IN **CYBER SECURITY**

INTRODUCTION TO THREAT INTELLIGENCE

Cyber security is a constant battle between attackers and target organizations seeking the upper hand. Security controls evolve to counter the latest attack methods and CKCs, while these methods evolve to outmaneuver these controls.

Threat intelligence is the result of analyzing data related to cyber criminals and their actions, understanding their motivation, targets, and behavior, encompassing past, present, and future activities. The data is analyzed using various tools and techniques to extract useful information about current and emerging threats relevant to the organization performing the threat intelligence analysis.

Consumers of threat intelligence can use the generated information to make informed decisions and take preventative actions such as adding or modifying security controls, changing behavior, or taking proactive steps.

Threat Intelligence provides organizations with prior knowledge of unknown vulnerabilities and possible attacks, offering a tactical advantage over adversaries – being forewarned is forearmed. The critical benefit of threat intelligence is it allows security defenses to be proactive and prevent attacks rather than reactive to limit the damage from attacks.

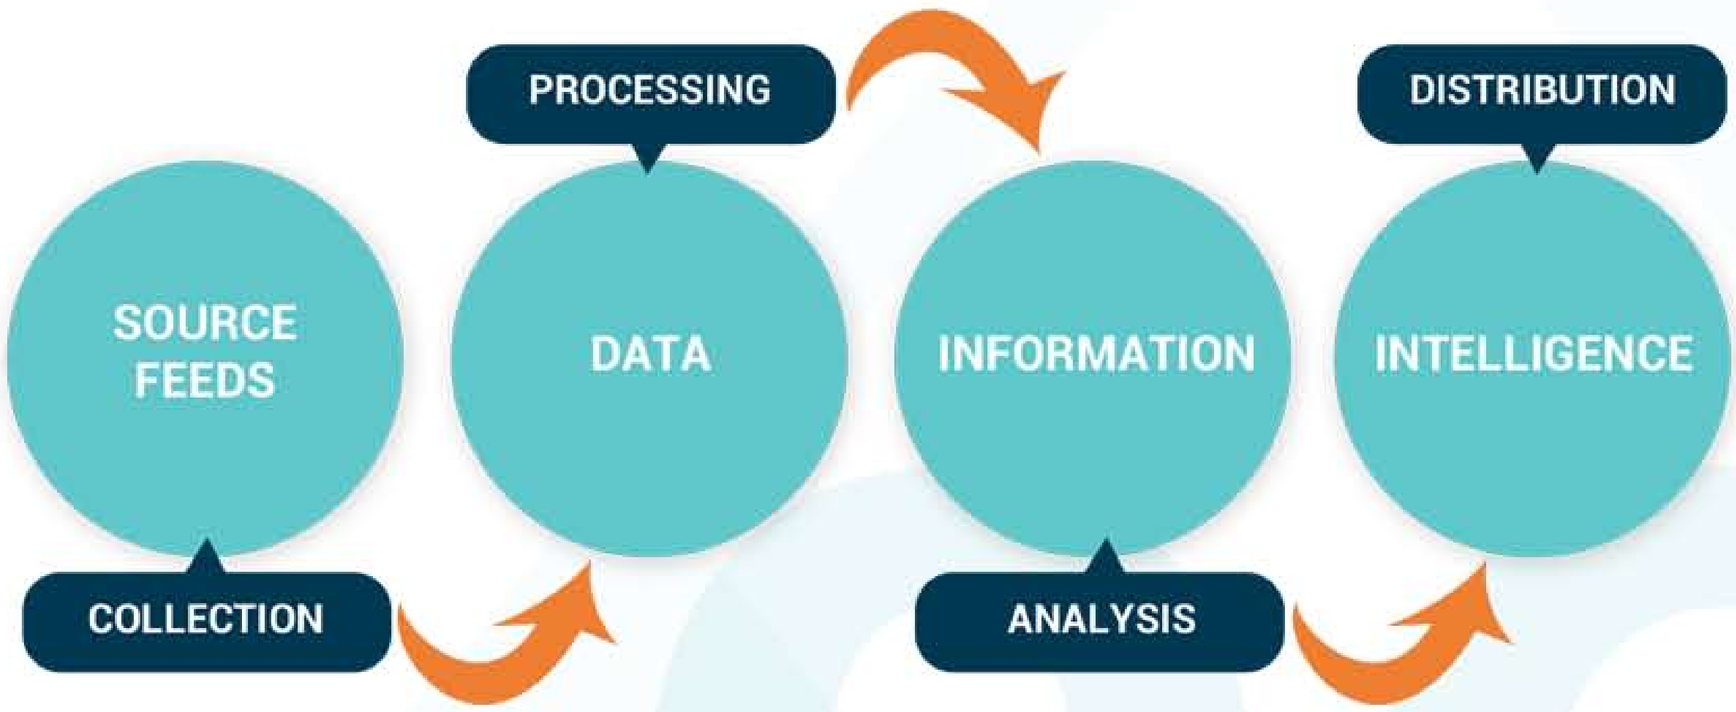

FIGURE 1 - Intelligence Generation Process

Data is the raw, unfiltered, unrefined information that comes in various forms, including communications traffic, textual symbols, vocal recordings, and media such as images. It can also include the output of systems and network sensors.

Information is data that has been processed, filtered, collated, and organized into a usable form ready for analysis.

Intelligence is the information that logical and analytical processes determine to be relevant, valuable, and actionable.

LMNTRIX

BE THE HUNTER | NOT THE PREY

TERMINOLOGY

An Advanced Persistent Threat (APT) is an attack played out without detection over a significant period. The capable attacker gains unauthorized access to a device, system, or network and remains undetected while collecting helpful information and conducting reconnaissance-type activities. The APT uses this time to study the compromised system to identify valuable data and opportunities to expand their access or increase their privileges. The ultimate goal of the APT is to gain access to all information of value and exfiltrate it into their systems or cause damage to the exploited system while having the ability to hide all traces of their activities.

The Cyber Kill Chain (CKC) is the sequence of actions comprising an attack from initial compromise to gain access through reconnaissance activities, malware introduction, information exfiltration, and lateral movement across or between different systems.

An Indicator of Compromise (IOC) is the evidence left behind by an attacker, such as malware files, records of logged events, and access to specific URLs and IP addresses. This information can be used to search for previously undetected attacks by tracing the presence of known IOC in system records and storage mediums.

THREAT KNOWLEDGE VS. INTELLIGENCE

Threat knowledge is the information about known cyber security events and incidents, ranging from the indicators of compromise from attacks that the organization has previously suffered to results of analysis of malware files that have infected a system.

This information forms data points of knowledge covering what is known about cyber security attacks and the perpetrators and enablers of these attacks. This information allows an understanding of what happened and why and can help protect against such attacks should they reoccur.

Typical threat knowledge includes details of the attacker, the type of attack, the method of delivery, the exploited vulnerabilities, the techniques used following the initial compromise, the impact of the attack, and the practices used to identify, respond, and recover from the attack.

The issue is that known attacks constantly evolve to defeat the protective measures created to prevent them, and new unknown attacks are developed to evade the latest protective measures.

Threat intelligence is the means of collecting and analyzing the threat knowledge to generate information that an organization can use to protect itself against not just known attacks but also novel attacks that can reasonably be expected to occur in the future. Threat intelligence fundamentally anticipates future threats based on historical threats, part projection from trend analysis, and part prediction.

Typical threat intelligence includes the target organizations by sector, location, or other discriminators, the motivations for attack, the technology that attacks targets, the nature of the exploited weaknesses or vulnerabilities, the delivery method, and the attack's objectives.

The challenge is that there are millions of threat knowledge data points, most irrelevant to one specific organization. The purpose of the threat intelligence process is to collect the data, extract the relevant information and perform an analysis that produces an actionable product.

THREAT INTELLIGENCE TYPES

Threat Intelligence can be partitioned into strategic, operational, tactical, and technical classifications.

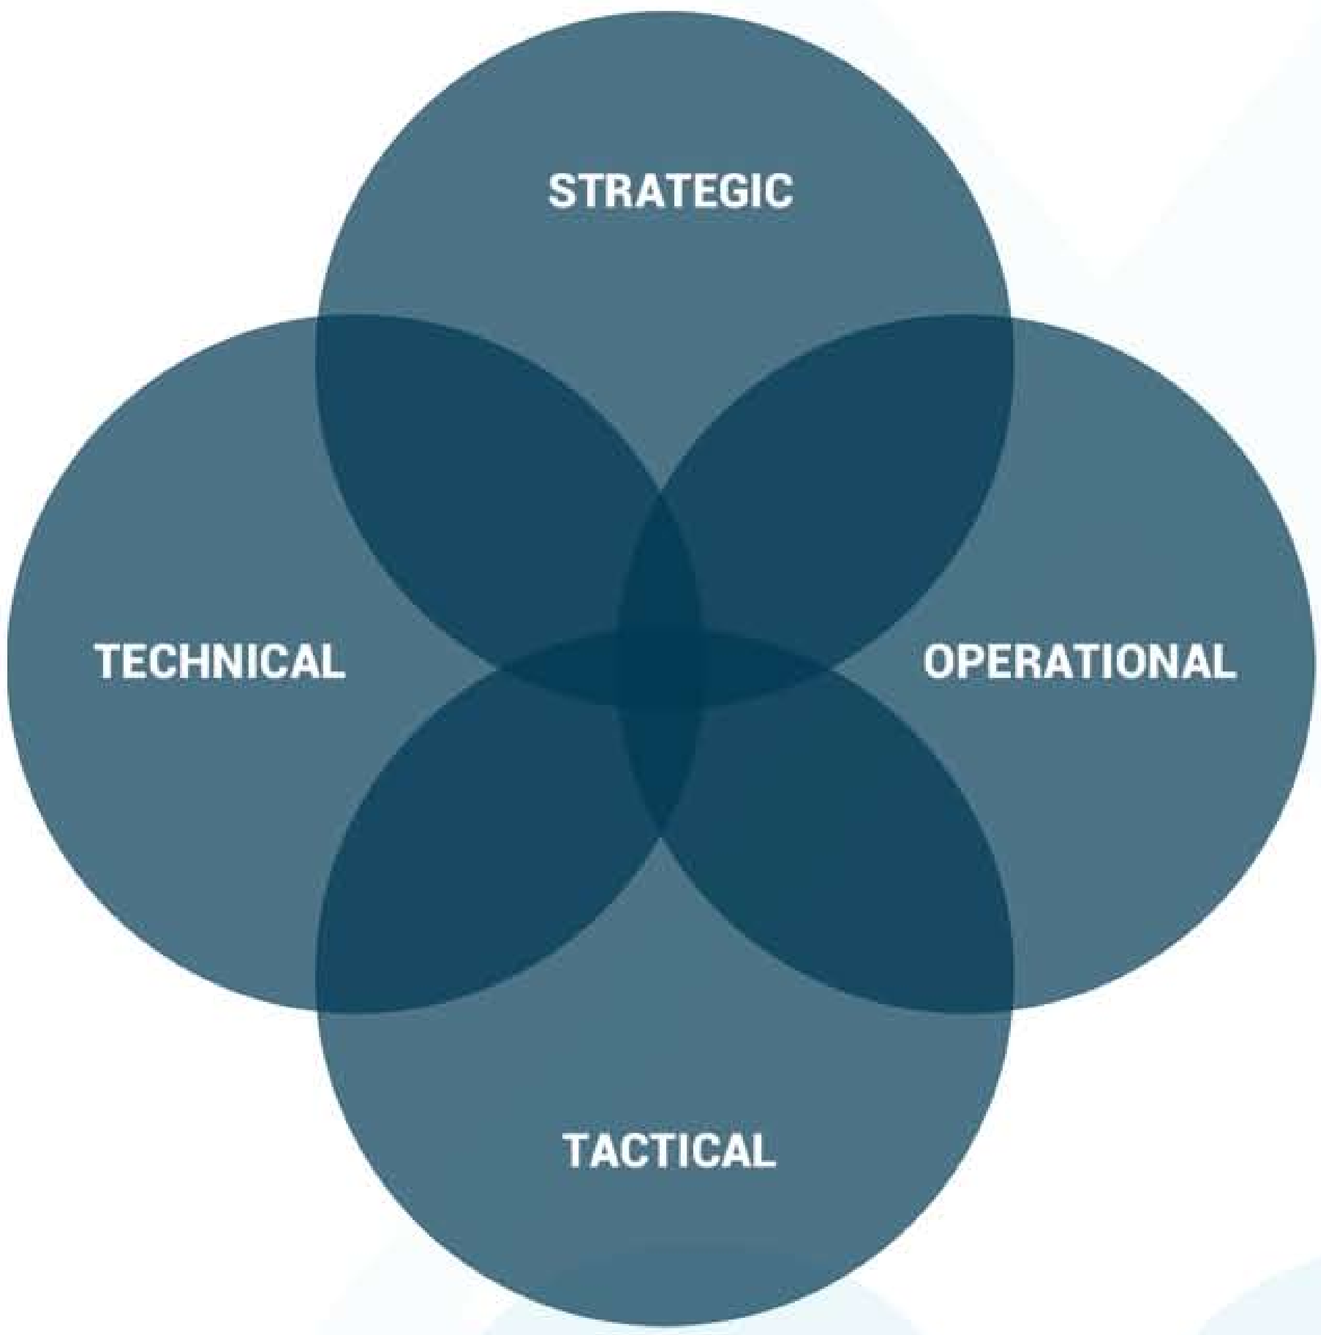

FIGURE 2 - Threat Intelligence Types

STRATEGIC THREAT INTELLIGENCE

Strategic threat intelligence offers an organization an overview of its threat landscape to inform long-term planning and decision-making. The threat landscape is the credible threats that an organization faces, determined by the operating environment, the nature of the organization's systems, the capabilities of its technical and procedural security controls, and the abilities of its security professionals. The better the systems, controls, and people, the smaller the threat landscape.

Strategic threat intelligence is often driven by economic and geopolitical factors, with world events influencing which hostile nation-states will seek to attack what type and location of a target and for what end goal.

At the strategic level, threat intelligence will provide insights into exposed vulnerabilities, security control weaknesses, and behavioral risks that could be exploited by an attacker using current or predicted attack vectors. Additionally, intelligence should provide indications of the likelihood and severity of potential attacks to support the prioritization of risk mitigation strategies and preventive actions.

Strategic threat intelligence requires the most significant effort to produce in terms of data collection and analysis activities.

OPERATIONAL THREAT INTELLIGENCE

Operational threat intelligence offers an organization an assessment of the immediate threat landscape to inform short-term planning and decision-making.

Operational threat intelligence allows organizations to profile attackers to understand their methods, motivations, and behavior better. It also enables tracking active campaigns to gain awareness of the organizations that will be targeted.

Operational-level threat intelligence will provide information on known attacks to which the organization may be vulnerable and specific targeted threats against the organization and its business sector. Threat information will include details of attack vectors and timing, attacker motivation, and exploitation routes obtained from intercepting or eavesdropping on attack planning discussions combined with indicators of compromise.

Typical operational threat intelligence will have a reasonable period of relevance before the information becomes obsolete. It is based on attacker behavior rather than specific attack types, but production requires significant analysis effort.

TACTICAL THREAT INTELLIGENCE

Tactical threat intelligence offers organizations a clearer picture of the current threat landscape, which allows the identification of immediate credible risks and indicators of compromise.

Tactical threat intelligence should provide actionable information, including details of specific attackers and their tactics, techniques, and procedures (TTP) that pose a credible risk to the organization.

Tactical-level threat intelligence will provide details of known exploitable vulnerabilities, the indicators of attack, and the measures available to reduce or eliminate the risk of exploitation. This has a dual purpose of allowing faster detection and identification of an in-progress attack and practical steps to prevent known attacks.

Examples of tactical threat intelligence include signatures of newly uncovered malware executables. Typical tactical threat intelligence will have a limited period of relevance before the information becomes outdated as attackers adjust to stay ahead to cyber control technology.

TECHNICAL THREAT INTELLIGENCE

Technical threat intelligence provides detailed information on known attack vectors based on the indicators of compromise for each reported security incident. Attack-specific information such as command and control IP addresses, compromised internet addresses, malware signatures, or phishing message content can be used to search for reuse in other previously unrecognized attacks.

Technical threat intelligence is dynamic in nature due to the evolution of attack vectors to prevent detection, with such information having a relatively short period of usefulness before it becomes outdated.

Examples of technical threat intelligence include specific IP addresses and URLs associated with attacks.

FIGURE 3 - LMNTRIX XDR – Intelligence Dashboard

THREAT INTELLIGENCE SOURCES

Threat Intelligence is collated from a broad range of different sources, known as feeds. These threat intelligence feeds are typically continuous streams of data relevant to an organization about attackers and their attacks.

Threat intelligence processes analyze these data feeds to produce actionable information for an organization.

Typical sources of threat intelligence data include:

- Digital forensic evidence from the organization's Indicators of Compromise (IoC) from previous at tacks and external sources.
- Malware Analysis from reverse engineering of malicious executables to understand the principles of operation, identifiable signatures, and details of any communications with external resources.
- Cyber Counterintelligence (CCI) sources that collect data from attackers misled into interacting with deceptive environments, including honeypots and malware sinkholes.
- Human Intelligence (HUMINT) sources provide information collected from attackers using infiltration to conduct eavesdropping or social engineering techniques to obtain information.
- Open Source Intelligence (OSINT) feeds provide threat-related data from open sources, including forums, message boards, intercepted messaging systems such as email and instant messaging, and search engines.

Threat intelligence typically is collated from thousands of different feeds, requiring bulk data to be collated, sifted, and analyzed to extract useful and, importantly, relevant threat information. This process involves interpreting raw data using security knowledge to determine the information's relevance, currency, and actionability for the consumer organization.

The extracted threat intelligence then needs to be delivered to the consumers of that information in an understandable and actionable form to be effective.

FIGURE 4 - Intelligence Sources

INTEGRATING THREAT SOURCES

Various types of threat information are available from diverse sources across the surface and dark web for multiple purposes in different formats. For example, an organization will have thousands of sources to choose from to collate feeds for its internal threat intelligence processes. Therefore, selecting the appropriate source for performing threat intelligence analysis is critical. A critical consideration for selecting information feeds is the data format and compatibility with the organization's aggregation processes and technology stack.

` Depending on vendor support, direct integration of threat data is typically available for security rated devices such as Endpoint Detection and Response (EDR) solutions, Intruder Detection Systems (IDS), Firewalls, and web gateways. Also, Security Information and Event Management (SIEM) or centralized log management system tools can be used to integrate data.

Data feeds can provide information in a standard format such as the STIIX or TAXII standards. Others may provide text that requires processing to convert into a consumable format, either as a data stream, standalone files, or posted to a webpage for access using a provided API or scraping tools. Other feeds may require manual processing before data can be used.

Threat intelligence providers who collect, aggregate, and process data from multiple sources and formats are also available. Typically, these providers combine original research with open sources of intelligence such as common vulnerabilities and exposures (CVE) databases, breach information, and vendor-supplied data. In addition, these intelligence suppliers typically provide API-based access making integration into analysis processes straightforward.

THREAT INTELLIGENCE PROCESS

The threat intelligence process starts with collecting unprocessed data from available sources and progresses to reporting actionable strategic, operational, tactical, and technical intelligence.

The threat intelligence process follows distinct lifecycle stages.

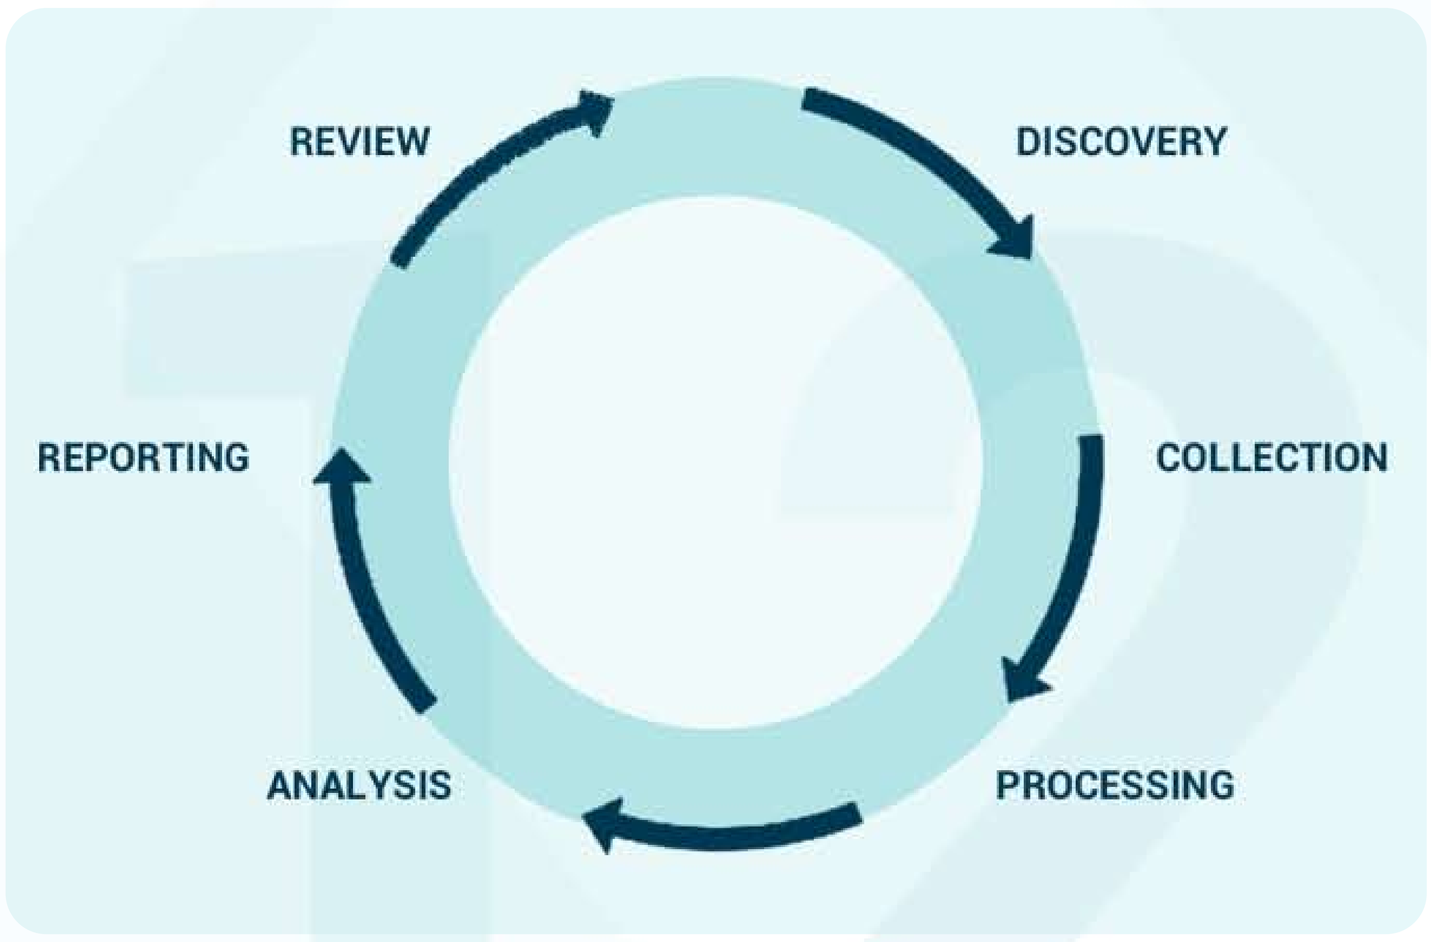

FIGURE 5
Threat Intelligence Process

DISCOVERY

The discovery phase defines the scope and sets the objectives of the threat analysis process. Understanding what business processes, systems, and assets require protection is critical to gather the correct intelligence.

The methodology for the threat intelligence process should be agreed upon and approved by stakeholders. The overall goals of the process are defined to support continuous monitoring of process effectiveness against key performance indicators and metrics reporting.

Collecting intelligence that is not relevant will reduce a business's security posture by consuming resources that should be expended on collating, processing, and analyzing relevant data.

COLLECTION

Threat intelligence sources relevant to the defined objectives provide raw data feeds that require collection into a data pool accessible for information processing activities. Threat intelligence can be collected from internal sources, including analysis of previous and current incidents. However, the most significant value comes from sharing intelligence between organizations to generate a complete picture of the threat landscape.

There are challenges with open intelligence sharing. Firstly, an organization may be reluctant to share intelligence about attacks they suffered to protect their reputation against that information becoming common knowledge with customers and suppliers. Secondly, intelligence information may contain sensitive or personal information covered by data protection laws and regulations governing how and when data can be legitimately shared.

Thousands of intelligence feeds are available from external sources, ranging from services provided by national security agencies, commercial security organizations, and open sources including social media, industry-specific agencies, subject matter experts, and private individuals.

Only reputable and relevant sources delivering data of acceptable quality must be included as intelligence sources. Low-quality data or feeds containing deliberately misleading information from disreputable or compromised sources can adversely impact processing by masking valid intelligence.

Low-quality sources may simply repackage other intelligence feeds, offering no new information but consuming data processing resources. Worst case, such feeds may amplify false or misleading intelligence and compromise the overall quality of the threat intelligence process.

Source quality can be measured by considering the number of unique indicators as a percentage of the provided indicators, the average age of Indicators of Compromise, and the number of reported false positives. Therefore, the threat intelligence program should continuously monitor source quality and positively bias information for high-quality feeds and downgrade information for low-quality feeds.

PROCESSING

Collated threat intelligence data will typically be delivered in diverse formats, structures, and content types. Intelligence will also often come from different countries using a range of languages, reflecting that for APTs, most attacks come from specific aggressive and hostile nation-states. Typically these are from those countries' government agencies, state-sponsored groups, or organized criminal collectives.

The processing of raw threat intelligence data is necessary to transform the data into a form that can be analyzed as a collated and coherent data set, mindful of the quality and reliability of the source of each data item.

Intelligent data processing techniques can accommodate the integrity and dependability of different intelligence feeds, so those from trusted sources can be given greater importance than those of lower quality, such as open-source feeds.

ANALYSIS

The analysis phase takes the processed threat intelligence data and extracts actionable information relevant to the organization. The information must be produced to address the organization's specific concerns and meet the goals agreed upon in the discovery phase.

REPORTING

Actionable threat intelligence information requires dissemination in forms that its audience can readily comprehend. Recipients will range from security teams responsible for investigating and resolving exploitable vulnerabilities to senior management with budgetary and task approval responsibilities. Each stakeholder for the threat intelligence process will have different requirements for the format and level of technical detail needed to perform their duties. A critical factor in the effectiveness of threat intelligence is the timeliness of reporting, reducing the window of opportunity for threats to be realized before the target organization has mitigated any risks that the threat poses.

REVIEW

All effective processes follow a recurring lifecycle to refine results and implement continuous improvement. Feedback from threat intelligence reporting allows adjustment of the lifecycle stages to improve the overall process's accuracy, relevancy, and timeliness.

The review feedback process also allows organizations to adjust their objectives and priorities to meet their business goals better.

LMNTRIX

BE THE HUNTER | NOT THE PREY

THREAT INTELLIGENCE STANDARDS

Threat intelligence relies on the automated sharing of intelligence, which necessitates the adoption of agreed data structures so information transferred between organizations can be correctly interpreted.

The MITRE Corporation has developed three standards for describing threat intelligence data and support sharing, STIX, CybOX, and TAXII.

- STIX is the Structured Threat Information Expression that shares vocabulary with CybOX for describing cyber threat information. It includes the definition of observables, indicators, incidents, tactics, technique and procedure (TTP), exploit targets, courses of action, campaigns, threat actors, and reports.
- CybOX is the Cyber Observable eXpression XML schema. This format characterizes the chronology and time range between events.
- TAXII is the Trusted Automated eXchange of Indicator Information, an open-source protocol and service specification for sharing actionable cyber threat information.

Other standards include:

- IODEF-SCI is the Incident Object Description and Exchange Format for Structured Cyber Security Information for normalizing data from various sources for human analysis and incident response.
- VERIS is the Vocabulary for Event Recording and Incident Sharing produced by Verizon for sharing incident data.
- OpenIOC is the Open Indicators of Compromise framework produced by Mandiant specifically for static information.
- RID is the Real Time Inter-Network Defense communications standard.

THREAT INTELLIGENCE PRINCIPLES

The Threat Intelligence process must adhere to the following principles to be fully effective, known by the CROSSCAT mnemonic.

- **Centralized** - Control of the Threat Intelligence processes should be centralized for efficient assignment of resources, correct prioritization, and effective communications management.
- **Responsive** - Threat Intelligence management processes should be responsive to the stakeholders' requirements consuming the information it produces.
- **Objective** – Threat Intelligence analysis processes should remain impassive and independent for management processes to ensure its product remains objective.
- **Systematic** – Threat Intelligence processes should follow methodical steps to deliver repeatable, reliable, and coherent information.
- **Sharing** – Threat Intelligence products should be distributed as widely as possible within any privacy constraints and protect vulnerable sources.
- **Continuous Review** – Threat Intelligence products should be continually assessed against new information to test validity, and results should be fed back into the Threat Intelligence processes for refinement and improvement.
- **Accessible** – Threat Intelligence products should be generated in forms that are accessible and understandable by their intended audience.
- **Timely** – Threat Intelligence must be delivered so it can be acted upon before it becomes outdated to have any value for the consumer.

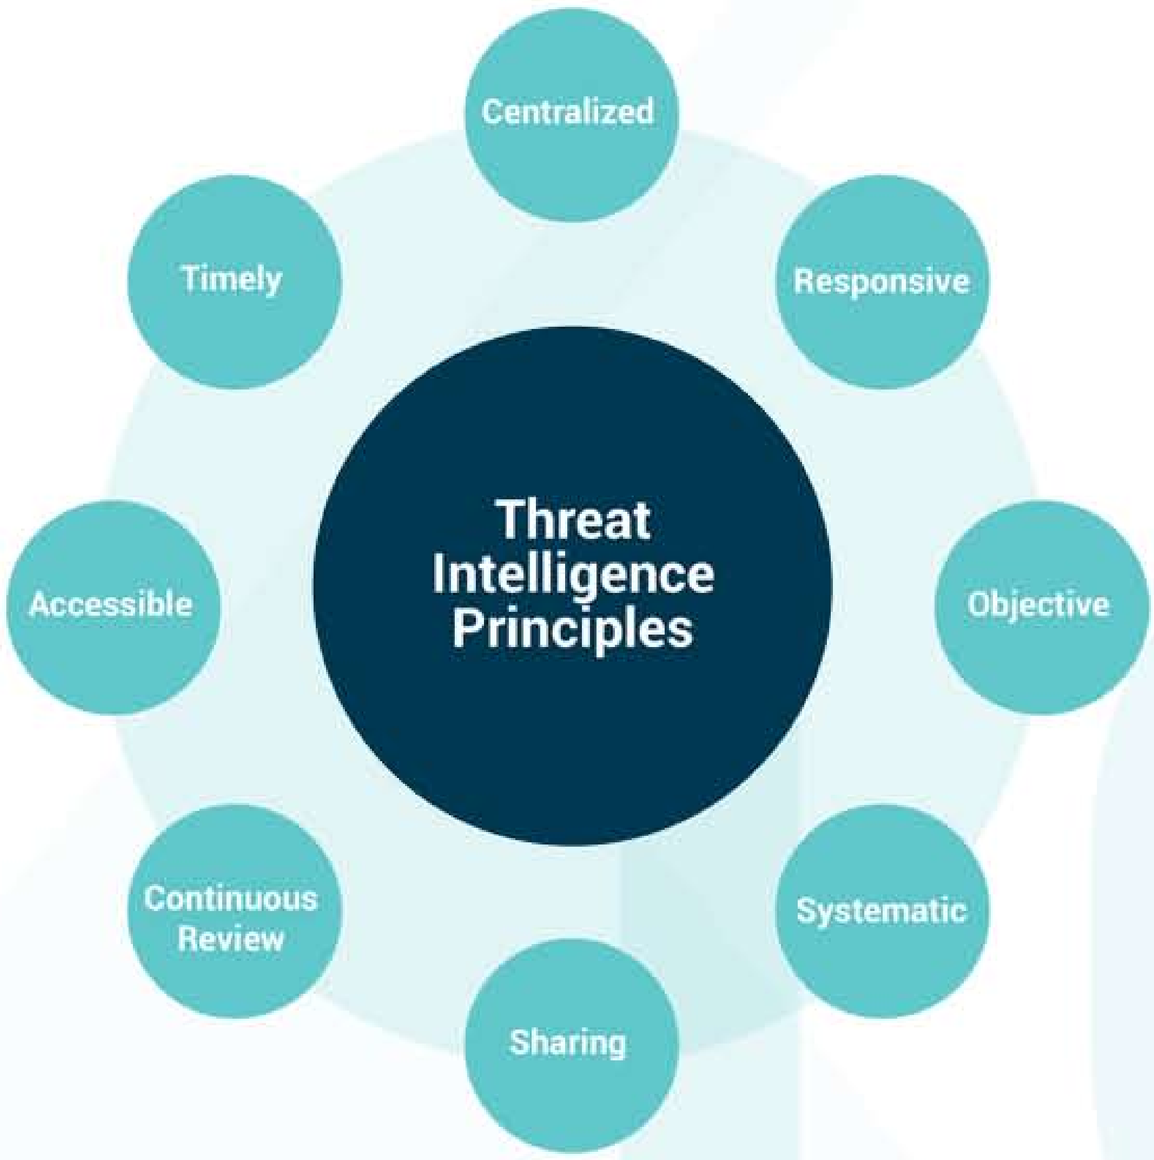

FIGURE 6 - The CROSSCAT Principles

LMNTRIX

BE THE HUNTER | NOT THE PREY

THREAT INTELLIGENCE BENEFITS

The intelligent application of threat intelligence into organizational security processes can significantly improve security posture and budgetary expenditure efficiency.

Threat intelligence offers organizations insight into novel threats before they are deployed by attackers, allowing preventative measures to be taken before an attack is launched.

- It offers organizations an insight into the motivations behind attacks, allowing them to modify behavior and influence decision-making to make the organization less likely to be targeted.
- It offers organizations information into attackers' TTP so that security controls can be assessed for effectiveness and additional countermeasures implemented as necessary.
- It offers organizations detailed information into known attack lifecycles and CKCs. Should an attack on the organization be identified, the security team can respond quickly and undertake informed recovery actions to minimize damage and disruption.

There is an added benefit for incident recovery in that the guidance provided by the threat intelligence allows the identification and response steps to be performed by fewer and less skilled team members than if they had no visibility of the attackers' TTP.

It allows organizations to more accurately quantify risks in terms of likelihood and severity to ensure the most significant threats are prioritized for mitigation measures and prevent investment in unwarranted protective measures that are not proportionate to actual risk levels.

It offers organizations information necessary to ensure security budgets are sufficient and target spending on the most critical areas of the business to mitigate the greatest threats to business operations.

UNDERGROUND **INTELLIGENCE**

THE DARK WEB EXPLAINED

The world wide web is the term used for the network of websites hosted on servers connected to the Internet. A necessary clarification is that the Internet is the name of the network infrastructure.

The world wide web is split into two parts: the part indexed by regular search engines, the surface web, and the other hidden part, the deep web. Websites that exist in the deep web are known as the dark web. These websites cannot be found using regular search engines and are not accessible using standard browsers. These dark websites require special tools and applications to discover and access their content.

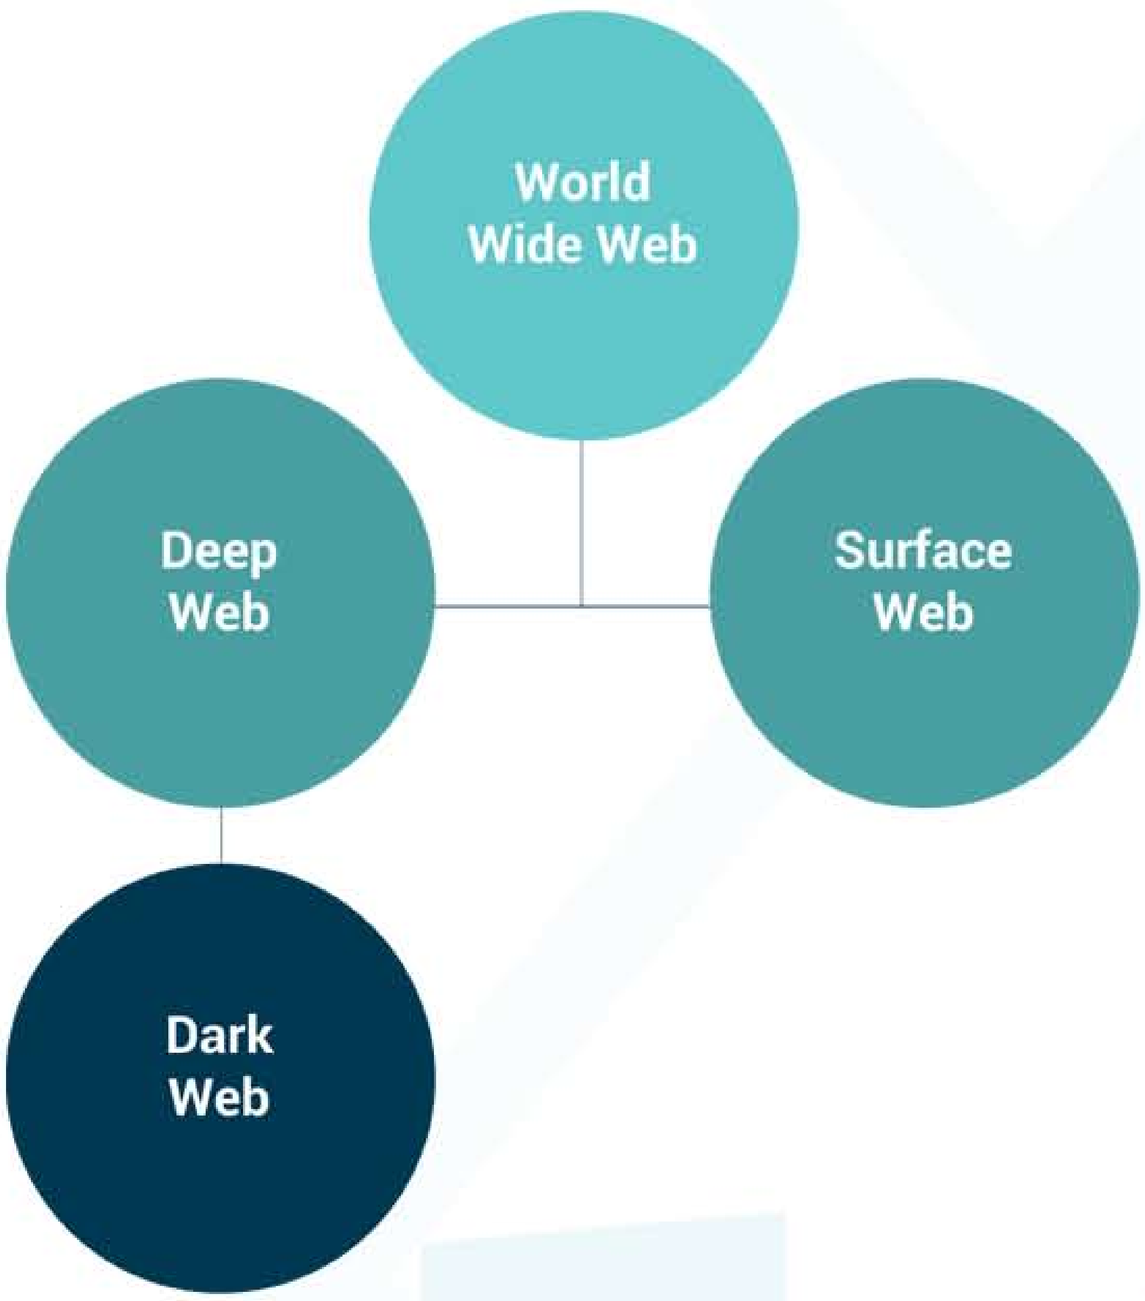

FIGURE 7 - Surface, Deep And Dark Webs

There are various reasons why websites exist on the deep web; the first is security reasons. The parts of websites that process payments or hold sensitive personal data will be protected using passwords and encryption so only authorized users can access information. These parts of the websites are on the deep web. Conversely, publicly accessible data such as product details in an e-commerce website will be on the surface web.

Another reason websites may be on the dark web is that their content is illegal, unethical, or immoral. These may include criminals selling stolen credit card information, illicit narcotics, or more disturbing materials. Consequently, the websites are hidden to prevent discovery by law enforcement agencies.

Alternatively, websites may be hidden to protect the privacy of visitors to the site. For example, a website publishing information shared between dissidents, human rights agencies, opposition supporters, and members of the public living in an oppressive dictatorship will need to protect the identity of visitors from internal security forces. Internet traffic to and from a website on the surface web can be easily tracked to monitor who uses the site from which locations. The dark web offers the ability to implement anonymous website access.

In the context of threat intelligence, the dark web contains websites, forums, and data storage used by attackers to advertise their services, discuss targets, and share their intelligence. They also host command and control applications, hold malware ready for deployment, and even hold copies of data stolen during successful attacks.

DARK WEB MONITORING

Dark web monitoring is the process of discovering and tracking information on the dark web. Typical applications include looking for data stolen from individuals or organizations, such as access credentials, usernames, passwords, or sensitive data, including network diagrams, firewall configurations, or other intellectual property.

A key benefit of dark web monitoring is that it provides evidence that an organization's systems have been breached if there are no visible indicators of compromise on the system itself. Analysis of the information uncovered on the dark web can indicate when data exfiltration is likely to have occurred and the extent of the breach. It may also show who performed the attack and the methods used to breach the organization's security controls.

Collating and processing data from dark web sources poses significant challenges for analysts due to the techniques criminals use to protect their identities and limit access to content.

- Criminals commonly use message encryption, steganography, and obfuscation techniques to prevent eavesdropping on their communications.
- Criminals often use multiple online identities and pseudonyms to make tracking their activities more difficult.
- Criminals will use web proxies or VPN connections, or anonymity tools such as the Onion Routing (TOR) facility to disguise or hide their geographic location. Additionally, a frequent change of routine will make tracking challenging.

However, with sufficient analytical processing resources, dark web monitoring can provide valuable intelligence into the plans and activities of cyber attackers by finding relevant resources and accessing their content. Attackers and their enablers often openly discuss vulnerabilities in operating systems, applications, and services. They will also share details of proof of concept and trade exploit code through dark web marketplaces. Attackers will also trade exfiltrated stolen data, including credentials, discussing availability, and trading data through the marketplaces.

INTEGRATING DARK WEB INTELLIGENCE INTO GENERAL THREAT INTELLIGENCE

The dark web offers organizations a valuable source of threat intelligence to supplement other data feeds. Monitoring the activities and conversations of attackers and the nature of exfiltrated data they publish can provide an informative insight into attack techniques, current targets, and upcoming plans.

The challenge for dark web monitoring is two-fold. Firstly, discovering intentionally hidden content requires a thorough understanding of the dark web and the tools available to access its content. The lack of any directory of dark web sites can make identifying credible sources of relevant and dependable threat intelligence on the dark web difficult. Search tools are available, but these will not provide complete visibility of all dark websites. Processes to identify website URLs from other intelligence sources offer the most reliable means of mapping out websites of interest.

Secondly, such sources of helpful information are buried among large volumes of other content you won't want to go near. Collating and filtering such content can require significant resources to achieve reliably in a sufficiently timely manner.

Typically, organizations will subscribe to third-party dark web monitoring services to resolve these challenges. This eliminates the effort required to map out and monitor intelligence sources, replacing this with a simple data feed that can be treated as other threat intelligence sources.

DARK WEB INTELLIGENCE TYPES

The dark web offers security analysts a range of different types of intelligence.

GENERAL INTELLIGENCE

Criminals and their enablers will often openly discuss the latest developments from new hacking tools, the latest vulnerability, or a new lucrative target for attack. The ability to monitor these conversations offers analysts important insight into the threat landscape and specific items of interest to the dark web community.

Attackers will often openly discuss a range of incriminating subject matters hidden behind perceived anonymity. However, detailed analysis of behavior, language, and subject matter can allow the compilation of sufficient data to create detailed profiles of individual attackers that can allow the deduction of their identities and tracking of their activities. Such intelligence can provide crucial information on attack techniques and target types to predict and prevent future attacks.

CRIMINAL SERVICES

Organized cyber criminals will offer services for purchase or rental as revenue streams for their group. These services can range from selling off-the-shelf malware ready for use or crafted phishing content prepared to be sent or hiring by the hour of networks of botnets able to perform a distributed denial of service attack. They can also include the purchase of access credentials for malware that they've silently deployed on compromised systems.

Botnets are collections of computing devices that a cyber-criminal can remotely control using malware installed in compromised devices. Internet of Things (IoT) devices are particularly vulnerable due to weak security controls and limited means to identify if the device has been compromised with malware. The devices that form the botnet can then perform automated attacks on the attacker's behalf, where a significant volume of simultaneous events is beneficial in overwhelming security controls. Typical botnets comprise many thousands of devices, and multiple botnets can be assigned the same task to increase the intensity of attacks if necessary.

Botnets offer criminals a lucrative income by renting their control via dark web marketplaces. Their attraction comes from their adaptability in being programmed to perform a range of actions with anonymity making tracing the source of any attack back to the person or group that hired the botnet difficult.

Monitoring what products and services are traded through the dark web provides valuable insight into attack methods, their CKCs, and the techniques, targets, and compromised systems. For organizations, this type of intelligence warns if they, specifically, or their industry sector generally, are in the attacker's sights. It also provides insights into previously unknown malware or indications of potential novel attack methods.

EXFILTRATED DATA

Following a successful attack that allowed attackers to exfiltrate information of value, the criminals responsible will often offer the exfiltrate data for purchase. This is particularly true for stolen access credentials. The attacker stealing usernames and passwords will sell these details for other attackers to use in their attacks rather than exploiting the credentials themselves.

Intelligence on the availability of compromised credentials provides insights into previous and current attacks on an organization. Often, the first IOC for an organization will be the appearance of their access credentials on dark web marketplaces. Also, individuals often learn that their personal details held by a commercial operator have been compromised by dark web monitoring services rather than from the breached business.

PHISHING ATTACKS

- The attacker sends a message masquerading as a legitimate instruction from a trusted party to initiate the attack.
- The recipient opens an attachment, clicks on a link or performs an action that starts the attack.
- The attack obtains sensitive information, installs malware or other malicious actions.

WATERING HOLE ATTACK

- The attacker compromises a popular online resource such as a social media site, ecommerce website or news publication website.
- Visitors to the compromised resource inadvertently download malware from the trusted website.

WEB ATTACK

- The attacker scans websites and other online resources for exploitable vulnerabilities.
- The attacker exploits found vulnerabilities and compromises the internet facing resource
- The attacker extends the attack across systems and networks to gain access to sensitive information or install malware.

MALVERTISING

- The attacker compromises online advertising to install malware.
- The compromised adverts are delivered to visitors of trusted websites.
- Visitors to the website who are presented with the compromised adverts unknowingly download malware.

FIGURE 8 - Common Attack Types

ACCESS CREDENTIALS

Access credentials for hacked accounts that can be exploited for gain will often be available for sale on dark web marketplaces. This includes e-commerce and financial management sites such as online banking, where exploitation can benefit the attacker financially. It also includes social media and email accounts that have value in launching phishing attacks on the contacts and connections of the compromised accounts.

Particularly valuable access credentials are corporate IT system accounts that allow attackers to gain a foothold from which to launch further attacks within the boundaries of the organization's systems.

Monitoring the dark web marketplaces can allow organizations to identify compromised accounts and investigate if they have been used to initiate response and recovery and, if not, take actions to prevent their use.

SYSTEM AND SOFTWARE VULNERABILITIES

An information type commonly traded between criminals is exploitable vulnerabilities in systems and software. Forums exist where attackers can discuss their efforts in hacking applications. At the same time, marketplaces allow the sale of everything from details of a valuable vulnerability to complete kits for compromising a flawed system.

When a software vendor issues a security patch, often, there is a concerted effort by analysts to reverse engineer the patch to uncover the vulnerability it resolves and the means to exploit it. However, once such information becomes available to attackers, there is a window of opportunity to take advantage of the flaw before the security patch is installed across the entire user base. Unfortunately, this window can potentially be many months or years long for organizations without a robust patching policy.

A previously unknown vulnerability in a popular application or operating system with proof of concept code that exploits the vulnerability with tangible results can command a very high monetary value for the person who discovers this zero-day CKC. Typically details will remain hidden while the exploit has value, though intelligence may garner insight into the type of attack and likely target software.

Monitoring forums and marketplaces provide intelligence into the vulnerabilities being discussed by attackers, the types of information being offered for sale, any breakthrough discoveries, and if vulnerabilities are being exploited in ongoing cyberattack campaigns.

TARGETS OF ATTACK

The information posted on dark web forums offers valuable intelligence into specific targets of attacks, where discussions include requests to attack a particular business or individual or attackers trading details of how to strike a specific target.

This information may include current or former employees seeking to damage a business as revenge against perceived injustice or disgruntled current or former employees offering information that could be useful to an attacker. Such discussions may also relate to a specific person, a former employer or partner, or a high-profile individual. There are numerous reasons why a person or business may become a target.

Monitoring dark web forums can provide advance warning that an organization may be targeted, allowing efforts to focus on uncovering the reasons behind the targeting and the likely attack methods that may need to be combated. In addition, timely intelligence offers the organization the chance to prevent targeted attacks by identifying and neutralizing the source of the threat.

UNDERGROUND INTELLIGENCE RESPONSE

One key piece of intelligence that dark web monitoring offers organizations is the ability to uncover previously unknown exploitation of their systems when their corporate information is discovered on the dark web. This enables the compromised organization to immediately act to mitigate the impact of criminals exploiting their systems and stealing their data. This timeliness is particularly critical where stolen data includes regulatory-controlled information or data owned by clients or suppliers.

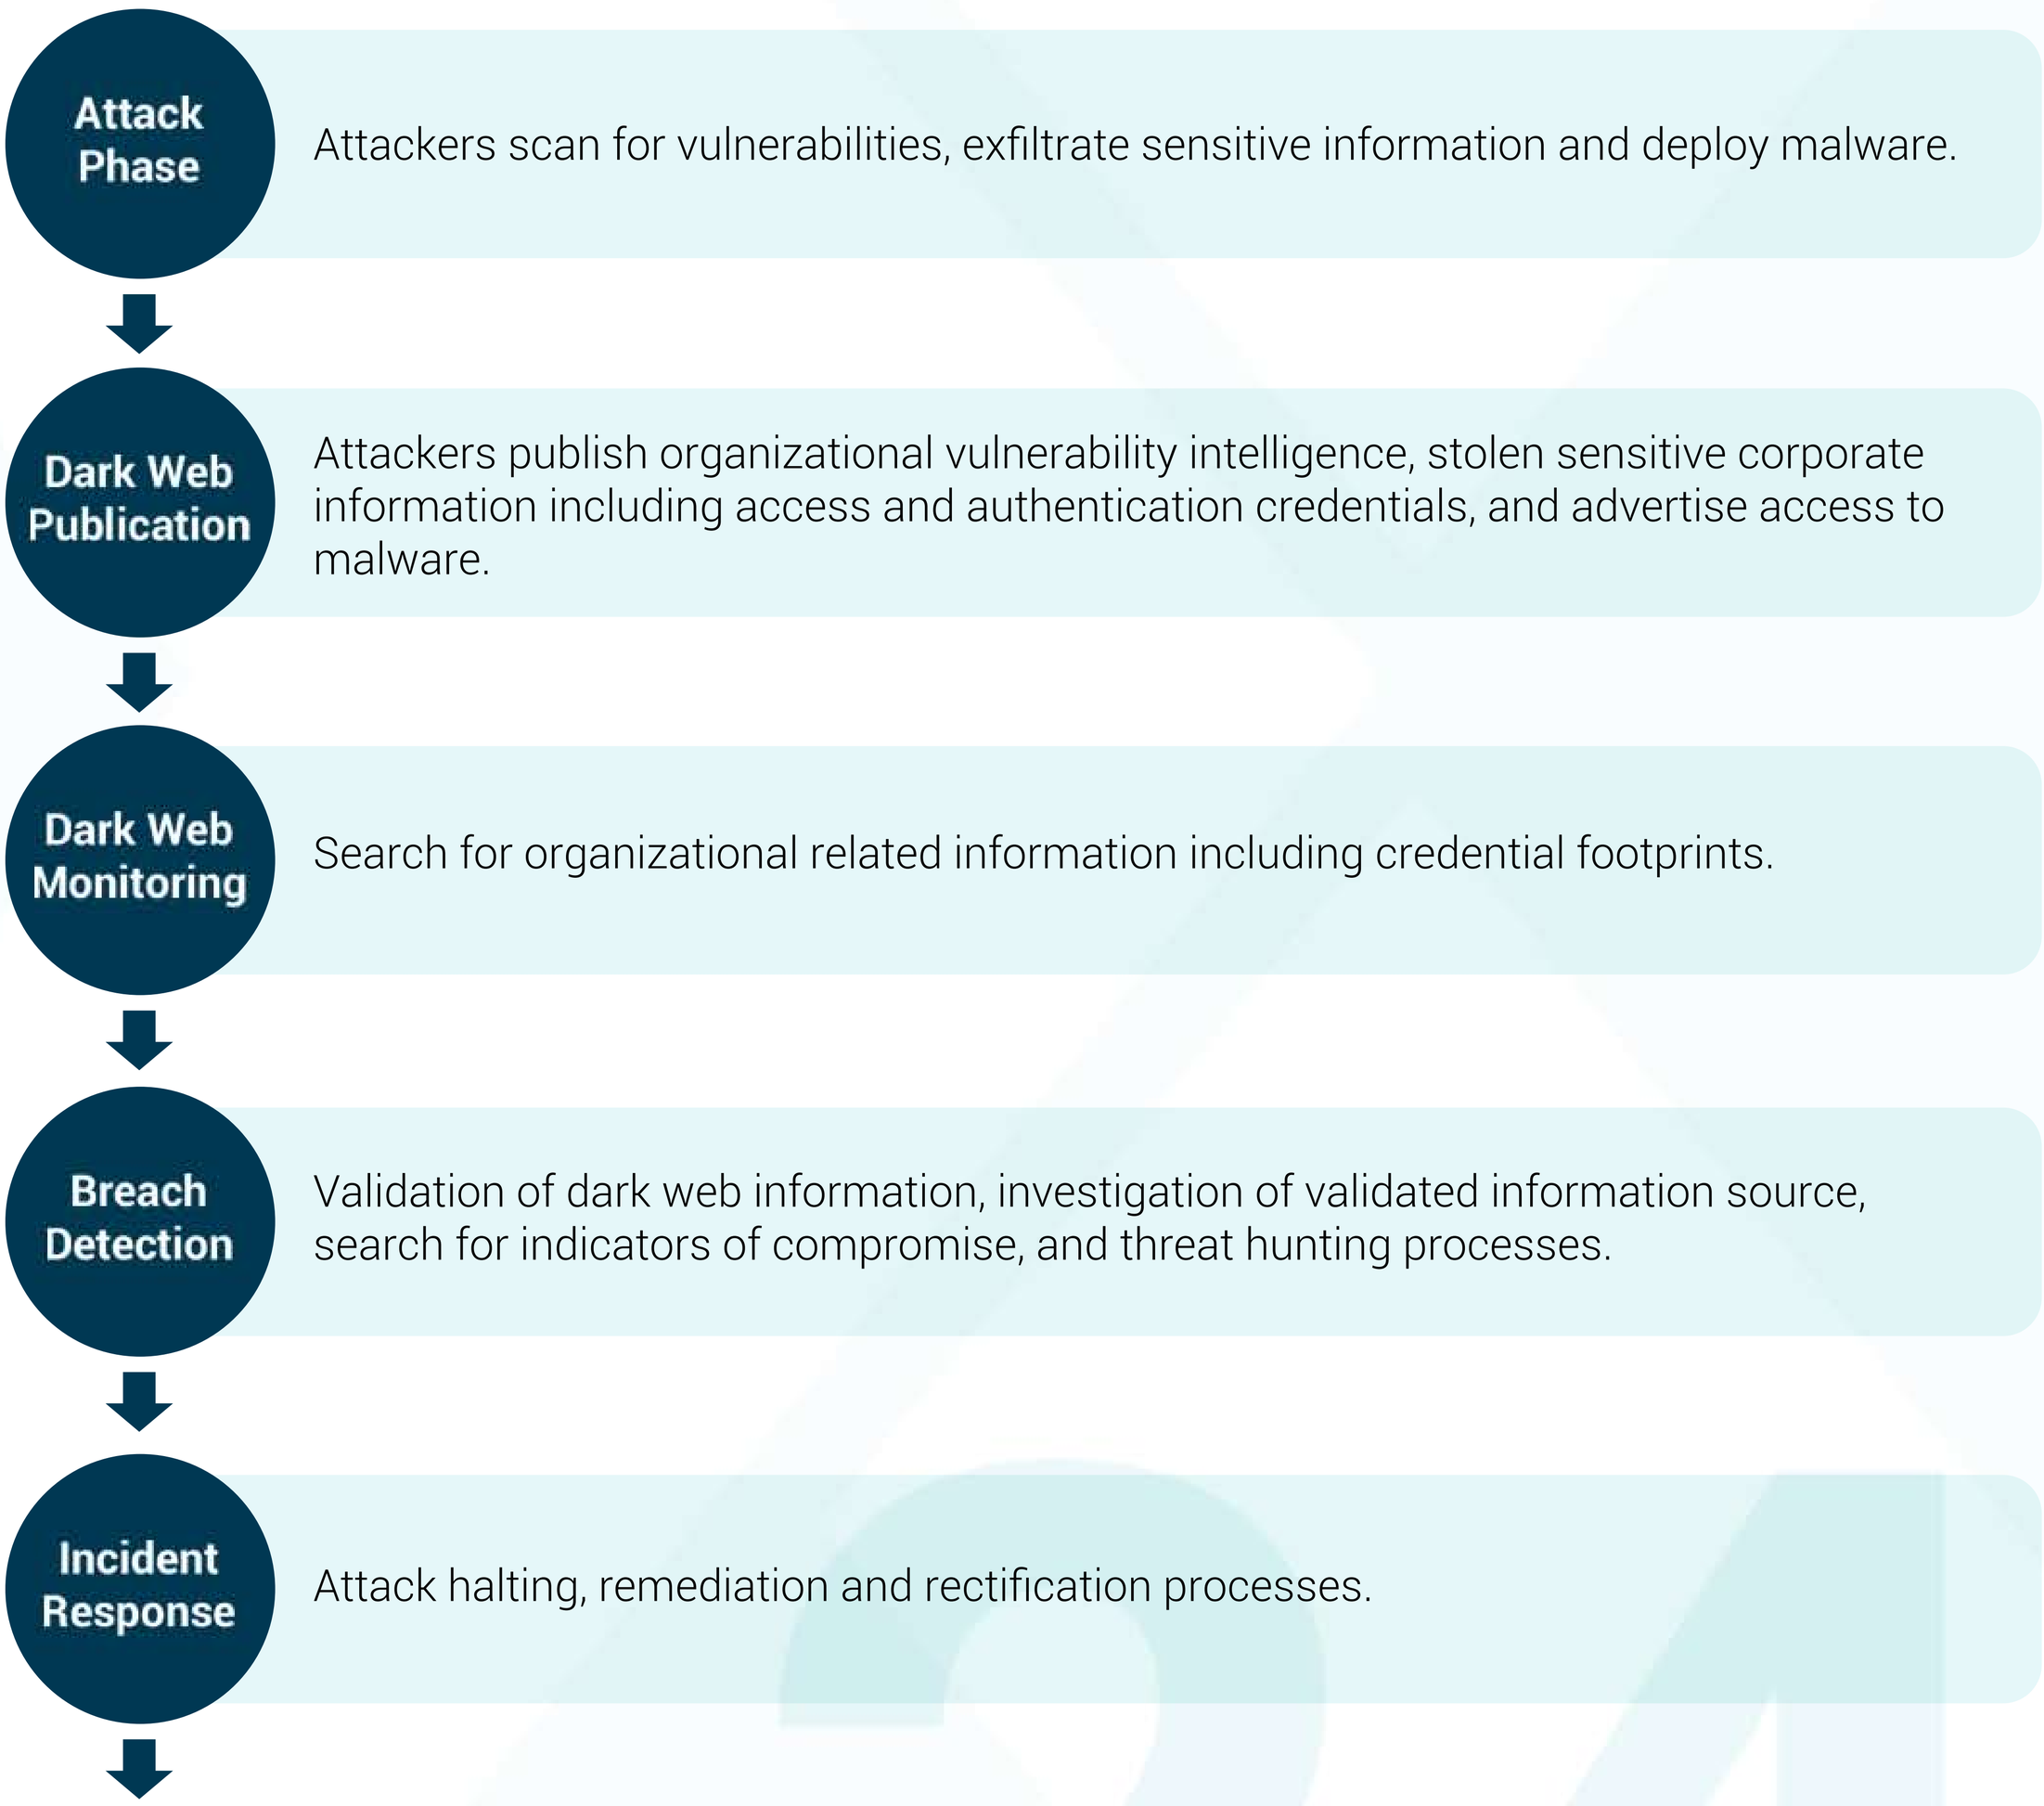

FIGURE 9 - Dark Web Response Process

The key measures that any organization should undertake in response to their data appearing on any part of the world wide web include:

CREDENTIAL MODIFICATION

All passwords, passcodes, and other access credentials should be changed to new and unique values, rendering any credentials stolen in a breach obsolete.

Where necessary, define and enforce a robust password policy that ensures new passwords cannot be easily guessed, defeated using brute force techniques, or derived from knowledge of any previous compromised values.

Multi-factor authentication (MFA) techniques should be applied wherever possible, using genuinely independent factors that a single-point failure in the authentication process cannot defeat. For example, the theft of a mobile device can give an attacker access to stored passwords and the ability to receive texted or emailed authentication codes for these passwords.

AFFECTED PARTY NOTIFICATION

All affected parties should be notified of the breach, including any clients and suppliers who may have sensitive information held on systems that were breached and system administrators of all external systems that have connectivity with the compromised systems. In the latter case, the attacker may have been able to move laterally across systems to compromise connections up and down networked supply chains.

MONITOR SYSTEM BEHAVIOR

All affected systems should be thoroughly investigated and monitored as part of the incident response and recovery actions to ensure that all attacks have been stopped and further exploitation prevented. This includes any third-party applications where access credentials to those applications may have been compromised.

SECURE RECOVERY PROCESS

Maintain sufficient system backups that allow recovery from any incident that results in loss of integrity or availability of corporate systems. This includes ransomware attacks that encrypt data, making it inaccessible without the decryption key. It also covers targeted attacks where sensitive information is exfiltrated and deleted from the breached organization's systems.

Backups should be protected to ensure any undetected compromise of a system does not compromise the integrity of the backup so that once the attack becomes active, recovery from the backup does not simply reinfect the systems or fail or restore systems.

MONITOR FOR CONSEQUENTIAL EFFECTS

Where an attack may have compromised access to systems and services such as banking services, e-commerce accounts, or social media services, these should be closely monitored for any unexpected or suspicious activity. This includes unusual operations, unauthorized changes, or other indicators of suspicious activity. In addition, any potentially compromised accounts should be frozen, and remedial actions taken to restore secure operation if necessary.

SECURITY IMPROVEMENTS

Evidence of a successful attack indicates deficiencies in the organization's security posture and should drive an improvement process. Therefore, this program should review the organization's technical, procedural, and personal security controls.

- Technical controls should be reviewed for weaknesses and vulnerabilities inherent in the control design or due to poor configuration and maintenance practices. If necessary technical controls should be replaced or augmented with additional controls of any type.
- Procedural controls should be audited to ensure their correct application and reviewed for any deficiencies in the protection they offer. Where possible, manually applied procedures should be automated or replaced with technical controls to reduce the risk of human error compromising organizational security.
- Personal security awareness should be reinforced using briefings and training packages to help improve threat recognition and reduce security risks associated with social engineering attacks, particularly phishing-type attacks.

LMNTRIX

BE THE HUNTER | NOT THE PREY

ACTIONABLE
THREAT INTELLIGENCE

USING THREAT INTELLIGENCE

How an organization consumes threat intelligence will depend on the nature of the information and the form in which it's produced. Actionable use of threat intelligence can be categorized into predict, prevent, detect, and respond actions.

- **Predict** - Strategic threat intelligence allows an organization to predict an attack and plan countermeasures using forecasts of evolving threat types and targeting information.
- **Prevent** - Detailed threat information, such as malware signatures, obtained from technical-level intelligence can allow an organization to prevent an attack by blocking malware at the perimeter of its systems.
- **Detect** - Technical and tactical threat intelligence provides information such as an attacker's tactics from their TTP, which allows organizations to detect threats within their networks and systems using threat-hunting processes.
- **Respond** - Detailed tactical threat intelligence also allows an organization to respond more effectively to attacks by using knowledge of the attacker's techniques and procedures from their TTP to inform how best to halt, remediate and recover.

THREAT INTELLIGENCE CHALLENGES

Threat intelligence feeds deliver vast quantities of unprocessed data about threats, attackers, incidents, and background noise. Some data will be relevant and current; however, most will be outdated or only applicable to specific geographical regions, industry sectors, target applications, operating systems, or other niches.

Actionable threat intelligence comes from filtering and analyzing raw threat data to produce relevant and current information that applies to the organization. Unfortunately, threat information feeds have significant levels of noise that can quickly swamp analysts with irrelevant or false information that hides the critical intelligence that can prevent an attack.

The currency requirement is critical as intelligence that becomes outdated before reaching security teams for action will waste vital resources hunting for a threat that has already passed or evolved beyond recognition using the available information.

The purpose of actionable threat intelligence is to empower security teams to identify, respond and mitigate credible security threats before the organization suffers significant harm or irreparable damage. To achieve this, security analysts need to recognize and prioritize credible threat information within the context of the organization's specific circumstances.

THREAT COMPLEXITY

As emerging threats evolve, they become more sophisticated, making detecting and responding promptly and effectively harder.

Significant advances have been made, including deepfake technology that can generate sufficiently believable content to improve spear phishing attacks' effectiveness. Increasingly techniques such as phishing are seen as an effective method of gaining access to systems, requiring less knowledge and experience to uncover and exploit vulnerabilities in technical perimeter controls without detection.

The increasing distribution of AI-enabled misinformation and the rise of disinformation-as-a-service is also being leveraged to facilitate social engineering attacks. Technical means of detecting fake media content are expected to maintain pace with the deepfake generators, but warning of an attack from threat intelligence will provide a more robust defense.

Ransomware attacks are becoming less effective thanks to compromised businesses' reluctance to pay any ransom, instead relying on system recovery techniques to restore services. This has seen a shift towards extortion-based attacks where attackers demand payment to stop sensitive information stolen in a data breach from being published in publicly accessible forums, relying on the threat of reputational damage to force payment.

A challenge for security processes is the sheer volume of vulnerabilities in a typical enterprise environment with new and legacy systems integrated across large and diverse user bases. This is compounded by the increase in the average number of endpoints, the continued rise in endpoint processing power and network speeds increasing the rate at which data is generated, and the relentless increase in attack types requiring the search for more potential indicators of compromise. This data volume will continue to grow, creating challenges for any analysis technology.

As a result, large organizations typically prioritize vulnerability management and patching processes based on assessments of the likelihood and impact of exploitation. Threat intelligence offers support by providing insight into the review of probability.

LMNTRIX

BE THE HUNTER | NOT THE PREY

THREAT DETECTION

Actionable threat Intelligence is critical for supporting threat detection and response processes as part of an organization's incident response practices.

Intelligence that provides TTP information for attackers is invaluable in proactive threat-hunting processes for detecting the presence of these attackers and their imitators. Additionally, understanding the attacker's motivations and abilities can ensure any response to their detected presence is effective, and the impact of the attack minimized. Finally, providing security teams with prior knowledge of attack vectors and their indicators with training tailored to their detection and remediation will maximize their effectiveness.

LMNTRIX

BE THE HUNTER | NOT THE PREY

MACHINE LEARNING
IN CYBER SECURITY

INTRODUCTION TO MACHINE LEARNING

Machine learning is the ability of a programmable device to learn an operation without being explicitly programmed to perform that operation. It's part of the broad field of artificial intelligence (AI), the ability to produce technology that exhibits intelligent behavior.

The purpose of employing AI-based systems is to perform complex tasks where systems cannot be programmed using formally defined processes due to time constraints or the inability to define the requirements deterministically. AI solutions that employ machine learning techniques achieve this by creating models from training data that allow the processing elements to recognize non-obvious patterns that can be applied to real-world data. This effectively allows the system to program itself based on learned experience.

The machine learning process requires large volumes of data to be collated and prepared to train the system. The greater the training data volume and quality, the more accurate the learned patterns and the better the machine learning outcome. The machine learning outcomes can also be manually refined when evaluation data does not produce the expected results.

MACHINE LEARNING BENEFITS

Machine learning systems can perform descriptive, predictive, and prescriptive functions. Descriptive functions use data to describe why an event occurred based on known experience. Predictive functions use data to predict what events may happen in the future based on behavior patterns. Finally, prescriptive functions use data to define what actions should be taken based on previous experience.

These descriptive, predictive, and prescriptive functions match the aims of applying threat intelligence to cyber security processes. They describe what attacks have occurred, predict what attacks are likely, and the actions to take for current and future attacks. Hence, machine learning is ideally suited for threat intelligence analysis processes.

MACHINE LEARNING PROCESSING TYPES

Three types of machine learning processes are available: these use supervised, unsupervised, and reinforcement techniques to train their models.

SUPERVISED MACHINE LEARNING

In supervised machine learning, the model is trained using data sets that have been labeled to indicate what the data means. The model will improve as it is exposed to more labeled data. This learning process type is commonly used for image recognition; pictures tagged with the subject matter allow the model to recognize the visual traits of each subject. The complete model aims to identify the subject in an unlabeled image to an acceptable accuracy level. Its benefit is automating manual processes that take understood data and follow proscribed steps.

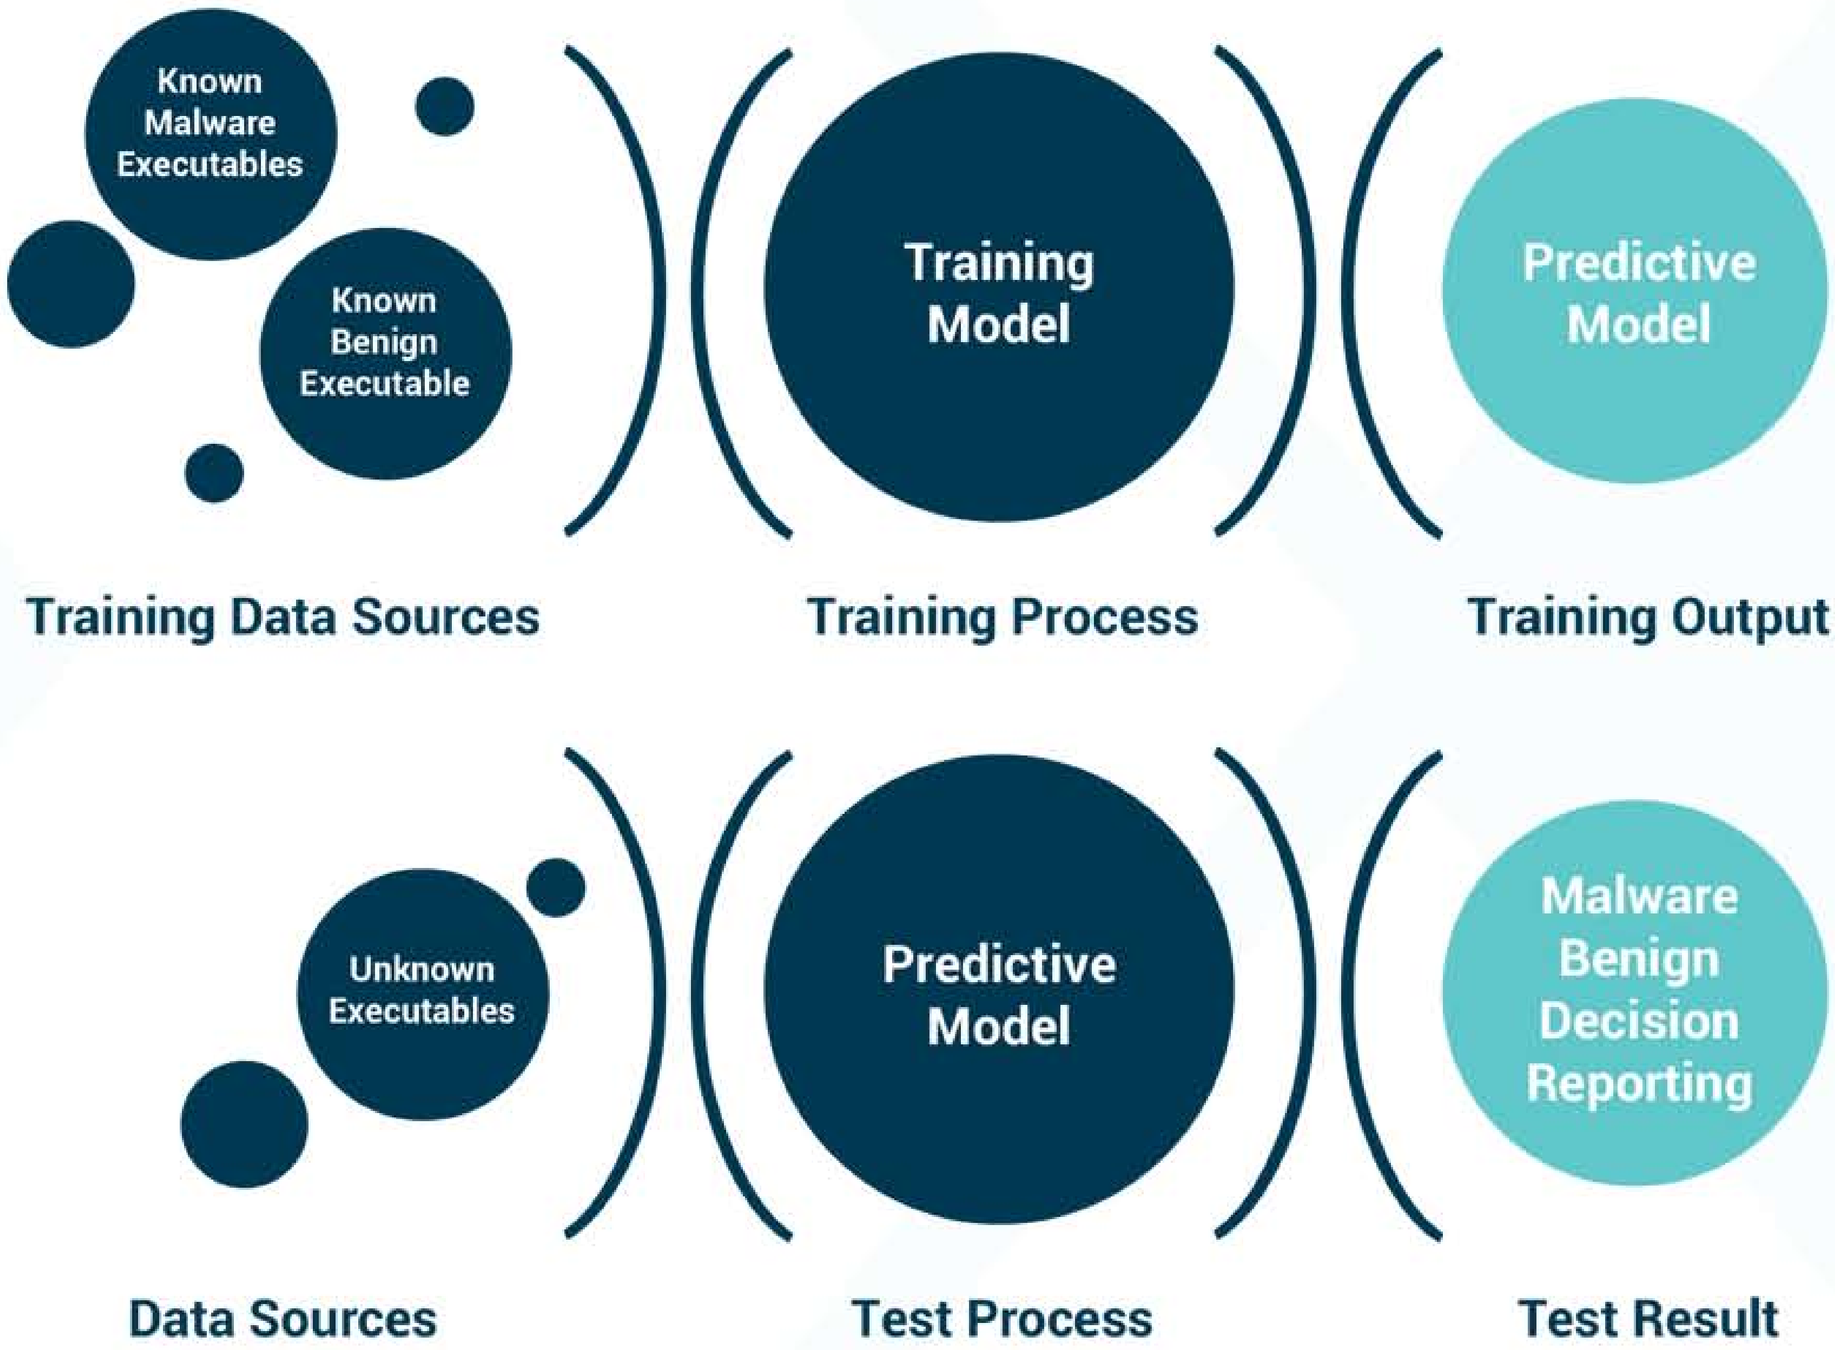

FIGURE 10 - Supervised Machine Learning Process

UNSUPERVISED MACHINE LEARNING

In unsupervised machine learning, the model is trained using unlabeled data where it is left to uncover non-obvious patterns or trends in the data that serve a useful purpose. This technique is useful in analyzing numerical data for forecasting-type applications. Its benefit is gaining insight from data where patterns are imperceptible.

BE THE HUNTER | NOT THE PREY

REINFORCEMENT MACHINE LEARNING

In reinforcement machine learning, the model is trained using a trial and error basis where desirable learned actions receive positive reinforcement to refine the model. This technique requires significant data volumes and time to produce meaningful results. However, its benefit is gaining insight or automating decision-making in processes that cannot be performed manually.

In addition to machine learning, applying AI processes to threat intelligence analysis requires natural language processing techniques. Therefore, the computer-based system is trained to understand natural language in its written and spoken forms for various languages. This allows the machine learning process to extract usable information from intelligence feeds that deliver textual or vocal content such as forum chat or eavesdropped telephone calls.

Another technique used in applying AI processes is using Computer Neural Networks (CNN). These replicate the principles of neural networks found in brains where large numbers of processing nodes are interconnected and arranged in layers to mimic natural processing methods. Labeled data is processed by each node performing a defined function to pass its results to connected network nodes. Creating multi-layered CNN allows for applying deep learning techniques where massive data volumes can be processed efficiently.

While deep learning techniques have the potential to maximize the extraction of threat intelligence from vast data sets collated from all available data feeds, they require significant processing resources that make them economically unviable for most organizations.

MACHINE LEARNING FOR THREAT DETECTION

Supervised machine learning techniques can detect attacks by analyzing security-related event information to identify patterns of potential abnormal activities that may indicate the presence of a malicious threat within the monitored system. The model can also include automatic response capability based on knowledge of known threat reactions and their effectiveness to learn which actions will provide the most beneficial response. In addition, threat information can be categorized based on attack type:

- The rate of Distributed Denial of Service (DDoS) attacks continues to increase year on year, expected to exceed 15 million in 2023, almost double the number experienced in 2018. Typical targets include gaming companies, broadband providers and hosting companies, cloud computing platforms, and cryptocurrency companies. Their increase in popularity is driven by the relatively low cost of botnet hire compared with ransomware attacks. Criminals can demand payment to halt attacks rather than expending effort bypassing security controls to deliver ransomware. They also offer nation-states options for disrupting the online services of other states using a sledgehammer approach to cause outages and infrastructure damage.

- Malware continues to threaten organizations, with, on average, five new zero-day vulnerabilities being observed each month, for which traditional signature-based anti-virus solutions will not be effective until the novel malware has been identified and analyzed.

- Ransomware remains a significant threat across all sectors and business types. However, the effectiveness of attacks is being diminished by increased awareness of recovery strategies that do not requirement payment of ransom demands.

- Threat actors are employing AI-based technology to improve the sophistication and effectiveness of their attacks. For example, automatically generated deepfake media files and coordinated release of disinformation can support social engineering and disrupt services.

- Social engineering attacks using Phishing techniques continue to be a popular attack vector due to the low investment and high returns these yield for attackers. Increased sophistication in attack methods and CKCs can result in the more prevalent use of more targeted attacks known as spear-phishing and whaling. Attacks are also moving away from email-based attacks to smishing and vishing attacks through mobile devices, which are harder to detect and prevent using technical security controls and for which awareness training is less mature.

- Supply chain targeting is gaining popularity as a means to breach robust perimeter controls of well-protected organizations by leveraging the access routes into their systems provided to less secure suppliers. As a result, third-party originating attacks rose from less than 1% of incidents in 2020 to a significant percentage the following year.

Data sets of known Indicators of Compromise provide training data supporting real-time detection and identification of threats based on attack behaviors.

MACHINE LEARNING FOR PHISHING PREVENTION

Technical controls for detecting and blocking phishing emails rely on techniques including recognizing malicious URLs or scanning for known content. Machine learning-based scanning offers predictive detection techniques capable of detecting patterns in email content indicative of phishing attacks that are effective even when URLs appear harmless.

As attackers move towards using SMS or voice messaging, such controls are more difficult to implement without the risk of disrupting normal business operations.

Supervised machine learning solutions are ideal for scanning message content. In the case of emails, this includes email headers and footers, language and terminology usage, patterns of punctuation, and other subtle distinguishing information to differentiate phishing emails from regular messages. In addition, for different messaging types, the adaptability of the machine learning solution allows them to become invisible protective barriers, including real-time monitoring of conversations to provide a warning or initiate in-depth monitoring if an indicator of a phishing attack is detected.

LMNTRIX

BE THE HUNTER | NOT THE PREY

MACHINE LEARNING FOR MALWARE DETECTION

Supervised machine learning techniques can detect malware by learning to recognize patterns in the executable code in known and labeled examples of benign and malicious executables to learn how to distinguish malicious code from ordinary software. Sufficient training can allow Machine Learning algorithms to identify novel malware that signature-based processes cannot detect. They achieve this through a more thorough analysis beyond instruction sequences to consider program behavior, API call sequences, network traffic, and structural entropy. This allows the model to distinguish between malware types such as trojans, viruses, ransomware, and adware.

There are two phases to malware detection. Primarily pre-execution analysis collects information about an executable file without performing its execution. Information from the binary data includes emulation results, file format and code descriptions, and other textual information extracted from the file. Additionally, post-execution analysis can provide behavioral information and details of any events that result from process activity.

A trained model can perform dependable predictive autonomous malware detection with sufficient training on an adequately sized data set. The accuracy of this model can benefit from additional supervised training when new malware variants are discovered from threat intelligence or malware detection in the wild.

The challenge with employing AI methods for malware detection is the importance of minimizing the false-positive rate that can disrupt normal business operations while managing the false-negative rate that can allow an attacker to gain a foothold in business systems.

Using a model trained in a representative environment that matches real-world conditions will reduce the risk of valid executables such as applications, utilities, and drivers being incorrectly labeled as malware.

MACHINE LEARNING FOR BEHAVIORAL MONITORING

Behavioral monitoring is the analytical and iterative investigation of past actions to generate insight into the validity of current measures using predictive techniques.

User Behavior Analytics (UBA) provides an additional layer to security monitoring by detecting atypical system usage patterns that may be symptomatic of abnormal system function due to a failure, a threat actor within the system, or malicious insider activity.

Behavioral analytics processes combine machine learning techniques with analytical data processing of users' behavioral data for any system or network. It can then create models of normal behavior to identify trends and patterns based on historical data and identify suspicious behavior that falls outside the expected parameters.

The machine learning model requires training using normal user behavior to learn what events fall outside those that constitute normal behavior to detect any abnormal activity. This technique is relatively straightforward, where users have strictly defined roles and perform consistently repeated actions. However, in the real world, where user responsibilities are complex and varied, change over time, and in response to specific events, collecting sufficient training data representing normal behavior can be a long and challenging process. Good quality and voluminous training data are vital if false-positive security events are to be minimized.

Any abnormal behavior that is sufficiently suspicious and triggers an alarm can then be investigated as a potential IOC.

Suspicious events may include accessing systems outside of regular hours, an unusually high number of accesses to a data store, or transferring exceptionally large volumes of data where such events are not normal behavior. However, they may be valid reasons for such behavior. Machine learning techniques can combine these prominent suspicious events with more subtle indicators from patterns of actions undertaken over significant periods to deliver more accurate results.

When applying machine learning to behavioral monitoring, the challenge is its limitations in detecting abnormal behavior in users such as developers, administrators, or other privileged users with irregular behavior patterns that do not fit regular usage models.

Behavioral monitoring is also better suited for sudden or dramatic changes in behavior associated with an active attack phase on an organization's systems. Any long-term sophisticated attack methods can fit within the acceptable tolerances of normal behavior and evade detection.

MACHINE LEARNING FOR THREAT INTELLIGENCE MONITORING / PROCESSING

The critical application of artificial intelligence solutions using machine learning techniques is processing large data volumes as typified by threat intelligence analysis. This technology is ideal for quickly handling massive data volumes to generate actionable results.

The key challenge with threat intelligence is the massive volumes of data available from sources that increase over time with no sign of plateauing in the perceivable future.

Employing deception technology and using techniques such as honeypots, sinkholes, and network sensors can dramatically increase the volume of data that requires processing.

Current threat intelligence analysis processes must handle more than a million threat data points in real-time and deliver accurate results. Using advanced automated solutions, collating and analyzing such massive data sets to detect anomalous events and behavior patterns is only achievable.

Machine learning offers the best solution for managing this data volume challenge without compromising the need for generating results sufficiently quickly for the results to be both actionable and ideally available before a threat materializes.

MACHINE LEARNING FOR THREAT RESPONSE

Machine learning-based processing of threat intelligence allows organizations to collate and process material relevant to any ongoing attack to respond quickly with the most effective response actions based on information gathered from previous responses to the same threat and knowledge of the TTP of the attacker.

Adaptive automated processes ensure the security team responsible for threat response receives accurate and timely advice, including solutions for halting ongoing attack and mitigating any subsequent actions by the attacker. While the critical benefit of automated response is speed, it offers additional benefits in reducing the risk of human error in response actions hampering an effective response.

Threat response time is a key performance indicator; sophisticated attacks offer short windows of opportunity for halting attacks between when an attack is initiated and becomes detectable through to when the attack is deployed, and the focus has to switch from damage prevention to damage limitation. Security teams can be significantly more effective when focusing on predicted attack prevention rather than ongoing attack reaction.

When an ongoing attack is detected, machine learning-based technology can gather data on the attackers' actions and consequential system effects for automated analysis to support response recommendations and decision-making. This technology can also include automated response actions to speed countermeasures and limit collateral damage.

- Machine Learning-based processes can halt the attack by automatically blocking the source of the event.
- Machine Learning-based processes can apply deception techniques to obtain additional intelligence about the attack techniques and the attacker performing the attack.
- Machine Learning-based processes can monitor the ongoing to obtain further information where intelligence value outweighs the need to halt the attack.
- Machine Learning-based processes can drive automated decision-making, including deciding to take no action if the event is deemed harmless or is not a threat to the organization.
- Machine Learning-based processes can assist in preventing the reoccurrence of this attack type by identifying and resolving exploited vulnerabilities.

The results from behavioral anomaly detection can use used to generate threat response recommendations based on learned experience to drive the response to an attack for specific policy-based attack scenarios.

MACHINE LEARNING FOR THREAT HUNTING

Incident response effectiveness during a cyberattack is significantly improved by taking active measures to find and identify new threats. Although, at the same time, threat intelligence can provide threat hunting with information on where and what to look for, machine learning technology allows organizations to improve the speed, thoroughness, and efficiency of threat-hunting processes.

Threat hunting is a vital technique for countering emerging threats not previously seen where their behavior is unknown, and no signatures or indicators of compromise are available to inform search processes. APTs able to loiter within an organization's systems collecting valuable information for extending the ongoing attack represent a significant risk if they cannot be detected before they move into a more active posture, exfiltrating sensitive information or inflicting harm.

Threat-hunting techniques are one of the few practical measures to counter silent attacks that focus on undetected data exfiltration for ransom-motivated attacks or to steal information or intellectual property to further their goals.

Zero-day exploits represent a persistent threat for security teams due to the challenge of detecting attacks without knowledge of the associated exploited vulnerabilities. However, those creating exploits typically rely on tactics and techniques used in previous attacks. Therefore, threat intelligence can be beneficial in providing information on attacker behavior that can indicate how new attacks will likely be undertaken that machine learning algorithms can use to identify potential indicators of novel attacks. In addition, the models of previous attacks can be used to highlight any correspondence between possible new incidents and known historical threats. The benefit of AI-based solutions using adaptive security models is the speed at which this can be achieved compared with typical manual threat-hunting processes.

Machine learning techniques can also support predictive forecasting based on historical data and its trends to forecast potential new threats before intelligence of their existence becomes available, countering threats before they are launched.

MACHINE LEARNING BENEFITS

The critical benefits of machine learning techniques in threat management include the following:

- The processing of vast volumes of threat data to extract relevant and actionable threat intelligence in a timely manner.
- The processing of increasing volumes of endpoint and network data for automated threat detection.
- Threat hunting for unknown attack vectors, including zero-day exploits and novel malware using predictive modeling techniques
- Automation of low-level, repetitive processes that enable human analysts to focus on the investigation of complex and nuanced indicators of potential compromise to deliver proactive threat detection

Machine learning technology is ideally suited for automating repetitive and time-consuming tasks associated with processing large datasets, such as threat intelligence analysis, searching data and network log files for abnormal events, monitoring network traffic for suspicious data, or scanning content for indicators malware.

Such tasks can be completed significantly faster while eliminating the potential for human error, typical with tedious and repetitive tasks that can result in a critical IOC being missed when buried amongst thousands of benign data points.

A related benefit is that automated processes can easily manage scaling challenges where a step jump in the data volume requires processing simply by provisioning additional processing resources. Conversely, scaling manual operations will require a significantly longer time and cost more.

Machine learning techniques can also be applied to an organization's systems, networks, and infrastructure to anticipate where weaknesses and vulnerabilities exist based on previous attacks to assess unmitigated risk levels. Combining risk with worst-case impact and exploitation allows the effective management of risk by prioritizing those areas where an attack is most likely to occur and has more severe consequences. This predictive risk assessment technique can maximize the expenditure of resources in improving security posture. Once additional security measures are introduced, the automated risk assessment process can be rerun to assess residual risk and provide a measure of the effectiveness of the measures introduced.

LMNTRIX

BE THE HUNTER | NOT THE PREY

AUGMENTED
THREAT INTELLIGENCE

AUGMENTED INTELLIGENCE

Augmented threat intelligence builds on the principle of augmented intelligence. Augmented intelligence is one element of artificial intelligence, operating collaboratively with manual processes to enhance human intelligence.

Augmented intelligence uses a combination of machine learning and deep learning techniques to process data and deduce actionable outcomes integrated into manual processes. Measures taken based on the results of the operations remain firmly a manual action. Augmented intelligence takes an advisory role to the human decision-maker.

Augmented threat intelligence integrates artificial intelligence-based solutions with human analysis to produce a partnership focusing on the human element. The goal is to boost manual cognitive performance and decision-making processes by leveraging the processing power of machine learning technology.

IMPLEMENTING AUGMENTED THREAT INTELLIGENCE

The augmented threat intelligence process automates collecting, processing, analyzing, and reporting functions, dealing with high volumes of complex statistical and narrative data. Machine-learning techniques support traditional rule-based, analytical, and statistical data processing to extract useful information from threat data. Typical functions implemented by augmented threat analysis processes include:

- Transform unstructured text using Natural Language Processing techniques that can handle information received in multiple languages to generate structured data suitable for processing.
- Extract structured knowledge from threat information based on connectivity between events, individuals, groups, locations, and other linking features.
- Find patterns and behaviors within data indicative of a novel threat, such as previously unknown malware.
- Forecast attack properties and events using predictive modeling driven by historical knowledge.

LMNTRIX

BE THE HUNTER | NOT THE PREY

DEPLOYING AUGMENTED THREAT INTELLIGENCE

Threat intelligence platforms will leverage multiple data sources to collect, process, analyze, and report information about security threats, vulnerabilities, and attacks. A comprehensive intelligence data set will require feeds from internal and external sources to be correlated into a cohesive data set. This involves intelligence aggregation from diverse sources, eliminating the need to manually aggregate and manage large amounts of threat-related data from thousands of distinct sources.

Consequently, an augmented threat intelligence platform will support faster and better-informed decision-making and enable a proactive security approach. As a result, offering organizations an off-the-shelf solution that helps security teams understand credible threats and potential risks.

The deployed augmented threat intelligence solution must fully integrate into the organization's security strategy to maximize benefits. Threat intelligence as a standalone solution will have limited effectiveness, relying on manual processes to match system events against received intelligence.

The key steps to deploy a fully integrated augmented threat intelligence solution include:

- The organization should ensure that the augmented threat intelligence platform consumes a comprehensive set of data sources that provide high-quality dynamic information relevant to the organization. This includes internal sources from any endpoint and network monitoring, behavioral monitoring, event data, outputs from security systems and solutions, and other indicators of compromise. In addition, external sources should be regularly reviewed for trustworthiness, quality, relevancy, and currency.

- Augmented threat intelligence processes must integrate with automated workflows to maintain manageable workloads for security teams. For example, threat data acquisition and processing for input to machine learning processes should be fully automated. Ideally, integration with autonomous incident management systems to allow cognitive-based prioritization driving automated alerting and remediation of high-risk incidents will reduce workload and enhance responsiveness.

- The organization should take a proactive approach to integrate threat intelligence into operational security policies and procedures. For example, intelligence can inform decision-making around defining access controls, permissions, and authentication processes to prevent attacks. Intelligence can also drive more efficient vulnerability patching processes with vulnerability identification and prioritization.

- Threat intelligence processes should integrate with existing security solutions such as Extended Detection and Response (XDR) or SIEM. The augmented intelligence processes can provide context for alerts and support autonomous detection and classification, allowing manual processes to focus on credible, high-priority, and fully documented events using an incident management solution.

- Threat intelligence processes should also integrate with incident detection and response processes to support better prioritization of which threats correspond to the most significant risk to business operations. Intelligence also allows better accuracy in predicting future actions an ongoing attack may take based on the attacker's previous behaviors and deduced intentions. Informed anticipation enables security teams to prevent further escalation of attacks by applying mitigating controls and initiating affirmative actions.

- Augmented threat intelligence needs to operate cooperatively with security teams and other stakeholders, providing clear information using data visualization techniques that enhance understanding and support the manual correlation of events. Threat information is only helpful when presented in an easy-to-consume format using methods such as dashboards, trend graphs, timelines, maps, and other visual presentations. Data type and representation should vary by audience, from high-level performance metrics for senior management to detailed reports for threat-hunting teams and security analysts.

The critical benefit of implementing a threat intelligence platform is the ability to deploy this solution as a software service or an appliance, physical or virtual, on-premises or in-cloud, private or public.

MEASURING AUGMENTED THREAT INTELLIGENCE EFFECTIVENESS

Implementing augmented threat intelligence will not realize its full benefit without a means to measure effectiveness. Poorly performing implementation will not be immediately apparent but will result in less than optimal security protection performance in the long term. This will often materialize as a greater-than-expected number of security incidents, though quantifying the expected number of incidents is itself challenging.

The general value threat intelligence delivers to an organization is understanding the current threat landscape to know what preventative measures must be taken and how to best prepare for the most probable potential incidents. The MITRE ATTACK framework is a helpful example of how threat intelligence can help an organization improve its security posture.

Key Performance Indicators (KPIs) can indicate the effectiveness of an augmented threat intelligence program. The critical factor is ensuring the selected KPI offers quality metrics relevant to the organization. Inappropriate KPIs can create biases that impede the implementation of the program.

- The main KPI for program effectiveness is the number of security incidents discovered using threat intelligence that would not have been detected in a timely manner without access to the information the program provided to the security team. This metric can also be extended to quantify improvements in the average time taken to detect and respond to incidents due to threat intelligence.

- One critical KPI relevant for all organizations is the false positive rate for reported attacks. A high false positive rate will cause the effort to be expended on nugatory tasks, depriving more essential functions of resources. This KPI demonstrates the effectiveness of the ability of the program to ignore irrelevant intelligence. Additionally, this KPI can provide positive feedback for the threat intelligence program by filtering those intelligence feeds that produce false positives via listing or cleaning mechanisms that bias those feeds that generate false positives.

- Another critical KPI relevant for all organizations is the false negative rate. Therefore, implementing an effective threat intelligence program should produce a demonstrable reduction in missed or late attack detection incidents.

- Another important KPI is alert prioritization accuracy. Low-priority alerts incorrectly categorized as high priority will result in analysts focusing on incorrect actions to the detriment of more important tasks. Similarly, high-priority alerts incorrectly classified as a low priority may be left unprocessed or delayed in processing. Correct categorization depends on the quality and reputation of the source of the threat intelligence, supported by the correlation of multiple disparate sources corroborating the threat's importance.

- A final important KPI is the quality and timeliness of reporting. Alerts should include sufficient relevant information to allow an analyst to conduct a thorough investigation and arrive at such time that the analysis can be completed before the threat intelligence becomes outdated.

BENEFITS OF AUGMENTED THREAT INTELLIGENCE

Machine learning techniques can be applied to augmented threat intelligence to support security incident analysis by automating repetitive tasks like viewing event log data, triaging voluminous alert information, or compiling reports. This frees the analysts' time to focus on analyzing high-risk events or conducting complex threat detection tasks that artificially intelligent systems may struggle to handle effectively.

The increasing sophistication of APTs, coupled with a greater incidence of low-level phishing and malware-type attacks, allows a focus on the higher-value decision-making associated with combating high-risk threats.

Human analysis is unlikely to be entirely replaced by artificial intelligence solutions. Still, machine learning allows organizations to focus on a reduced specialist team augmented with a machine learning-based security solution or maintain their existing team and use intelligent systems to provide more thorough and far-reaching protection.

The current trends around IoT technology adoption are also generating a significant increase in the number of endpoints requiring protection, compounded by the low maturity of security controls implemented in these IoT devices. More endpoints create more event data for analysis while offering additional vulnerabilities for exploitation, widening the attack surface. Automation of threat detection and response processes is essential in managing this exponential growth in the security management overhead.

USE CASES

USE CASE #1: ALERT HANDLING

Security Operations Center staff managing high-risk businesses such as financial service providers offering online services will typically need to handle large volumes of security alerts generated by the network traffic to the website. Alerts may be due to an actual attack or, more likely, to lower-risk scanning, probing, or user error.

Typically, analysts will spend time triaging each event to establish if it is a high-risk security incident, a low-risk event, or a zero-risk false positive alert. The triage process will entail an analyst cross-referencing the event against other sources of information to classify the event appropriately. Where analysts spend significant periods processing false positives, there is a considerable risk that this work's tedious and repetitive nature may mean they miss the subtle indications that an attack has generated an event. Additionally, humans are frequently fallible, particularly when placed under time pressure or demotivated by the nature of the task.

A sudden event increase can also be challenging to manage with manual analysis processes. For example, pulling in additional analysts may not be possible. In addition, requiring an existing team to work longer to process more events can exacerbate the risk of errors that miss an IOC.

Machine learning techniques are ideal for performing such repetitive and proscribed tasks automatically with the twofold benefit of eliminating human error and being capable of near-instantaneous scaling throughput to match the rate that events are generated. Rules-based processes for identifying abnormal and suspicious behavior are prime candidates for automation. Such processes can triage events and pass only those deemed high risk to the security analysts for further investigation, including providing details of all the supporting information to the suspicious event as a complete dataset to aid the manual analysis.

This approach will allow the analysts to focus entirely on more challenging investigative tasks, improving motivation and eliminating monotony. It also reduces the time spent on event analysis so the security team can spend more time on productive tasks, including proactive threat hunting.

USE CASE #2: NETWORK TRAFFIC MONITORING

Network traffic monitoring is an established method of threat detection, using techniques such as packet inspection and behavioral monitoring to detect suspicious traffic that could indicate an attacker's presence.

Such traditional traffic monitoring techniques are challenging to apply to more complex traffic types, such as File Transfer Protocol (FTP) traffic associated with application developers accessing centralized code repositories as part of their regular work duties. For example, frequent but irregular down and uploading of executable code across the corporate network by collaborative teams of developers can create challenges when determining the difference between ordinary and suspicious activity.

Individual developers will have different patterns of behavior that will vary between the different phases of the development lifecycle. For example, some developers may focus on working with a single file, while others may access dozens of separate files, such as library functions. In addition, developers with a quality control function may access hundreds of files.

Determining the static rules defining "normal" behavior to inform behavioral monitoring algorithms would be nearly impossible. However, machine learning processes could use the continuous monitoring of usage patterns to create a model of normal user behavior for every individual user to sufficient accuracy that alerts for suspicious behavior could be generated with adequate confidence in capturing attack indicators with a low false positive rate. In addition, correlating anomalous user behavior with other attack indicators from alternate intelligence sources can provide these required accuracy levels.

More importantly, such automated model-based processes can be maintained efficiently to manage changes to individual roles and responsibilities and the onboarding of new staff in a dynamic work environment.

USE CASE #3: AUTOMATED THREAT HUNTING

An organization looking to increase proactive threat hunting to counter growing threats to the business typically faces a resourcing challenge. Traditionally threat hunting required a team of knowledgeable and experienced security analysts supported by threat intelligence to inform their processes. However, as threats evolve, novel threats appear, and new vulnerabilities are found, these analysts would require training and potentially access to new tools to undertake their activities.

The combination of changing threats and vulnerabilities creates a significant effort for the analysts to manage. For example, a new vulnerability identified in an endpoint requires the team to review the library of existing threats to identify any that may have exploited the vulnerability and add additional considerations in the hunt for novel threats to that endpoint. Typically this involves the analysts collating and assessing data from CVE databases, endpoint security scans, system logs, and other relevant security solutions.

The results of this process inform the automated playbooks employed in the threat detection and response processes. Typically this would take several weeks to generate the playbook, including verification and review processes.

Machine learning-based processes can replace this manual procedure by fully automating the threat assessment and playbook creation to turn weeks into hours, significantly restricting the window of opportunity for an attacker to exploit a newly uncovered vulnerability before the risk is mitigated.

LMNTRIX

BE THE HUNTER | NOT THE PREY

HOW **LMNTRIX XDR** BENEFITS FROM **AUGMENTED THREAT INTELLIGENCE**

LMNTRIX ACTIVE DEFENCE

LMNTRIX Active Defense is a three-tier outcome-based XDR solution that is an enhanced development of the traditional Managed Detection and Response (MDR) service. It is composed of the following tier elements:

- **LMNTRIX XDR** (AWS Data Lake and Platform)
- **LMNTRIX Technology Stack** (Deployed deep within Customer Networks)
- **LMNTRIX Cyber Defense Centre** (Security Analyst Driven).

It is supported by specialist threat intelligence services, including **LMNTRIX INTELLIGENCE** and **LMNTRIX RECON**.

- LMNTRIX Intelligence
- LMNTRIX Recon

FIGURE 11 - LMNTRIX Active Defense

LMNTRIX

BE THE HUNTER | NOT THE PREY

LMNTRIX XDR

LMNTRIX XDR natively unifies Machine and Underground Intelligence, NGAV, EDR, NDR, UEBA, and Deception Everywhere with completely automated attack validation, investigation, containment, and remediation on a single, intuitive platform. Backed by a 24/7 Managed Detection and Response service at no extra cost, LMNTRIX provides comprehensive protection of the environment for even the smallest security teams. In addition, it is a single investigative platform for insights into threats on enterprise, cloud, hybrid, and industrial control systems (ICS) networks.

The **LMNTRIX XDR** delivers unique advantages over current network security solutions. It is a holistic and multi-vector platform with an unlimited retention window of full-fidelity network traffic, innovative security visualizations, and the ease and cost-savings of an on-demand deployment model.

LMNTRIX XDR is based on multiple detective, responsive, and predictive capabilities that integrate and share information to build a security protection system that is more adaptive and intelligent than any one element. The constant exchange of intelligence between the Active Defense components and the broader cybersecurity community enables LMNTRIX to keep abreast of the TTP of the most persistent, well-resourced, and skilled attack groups.

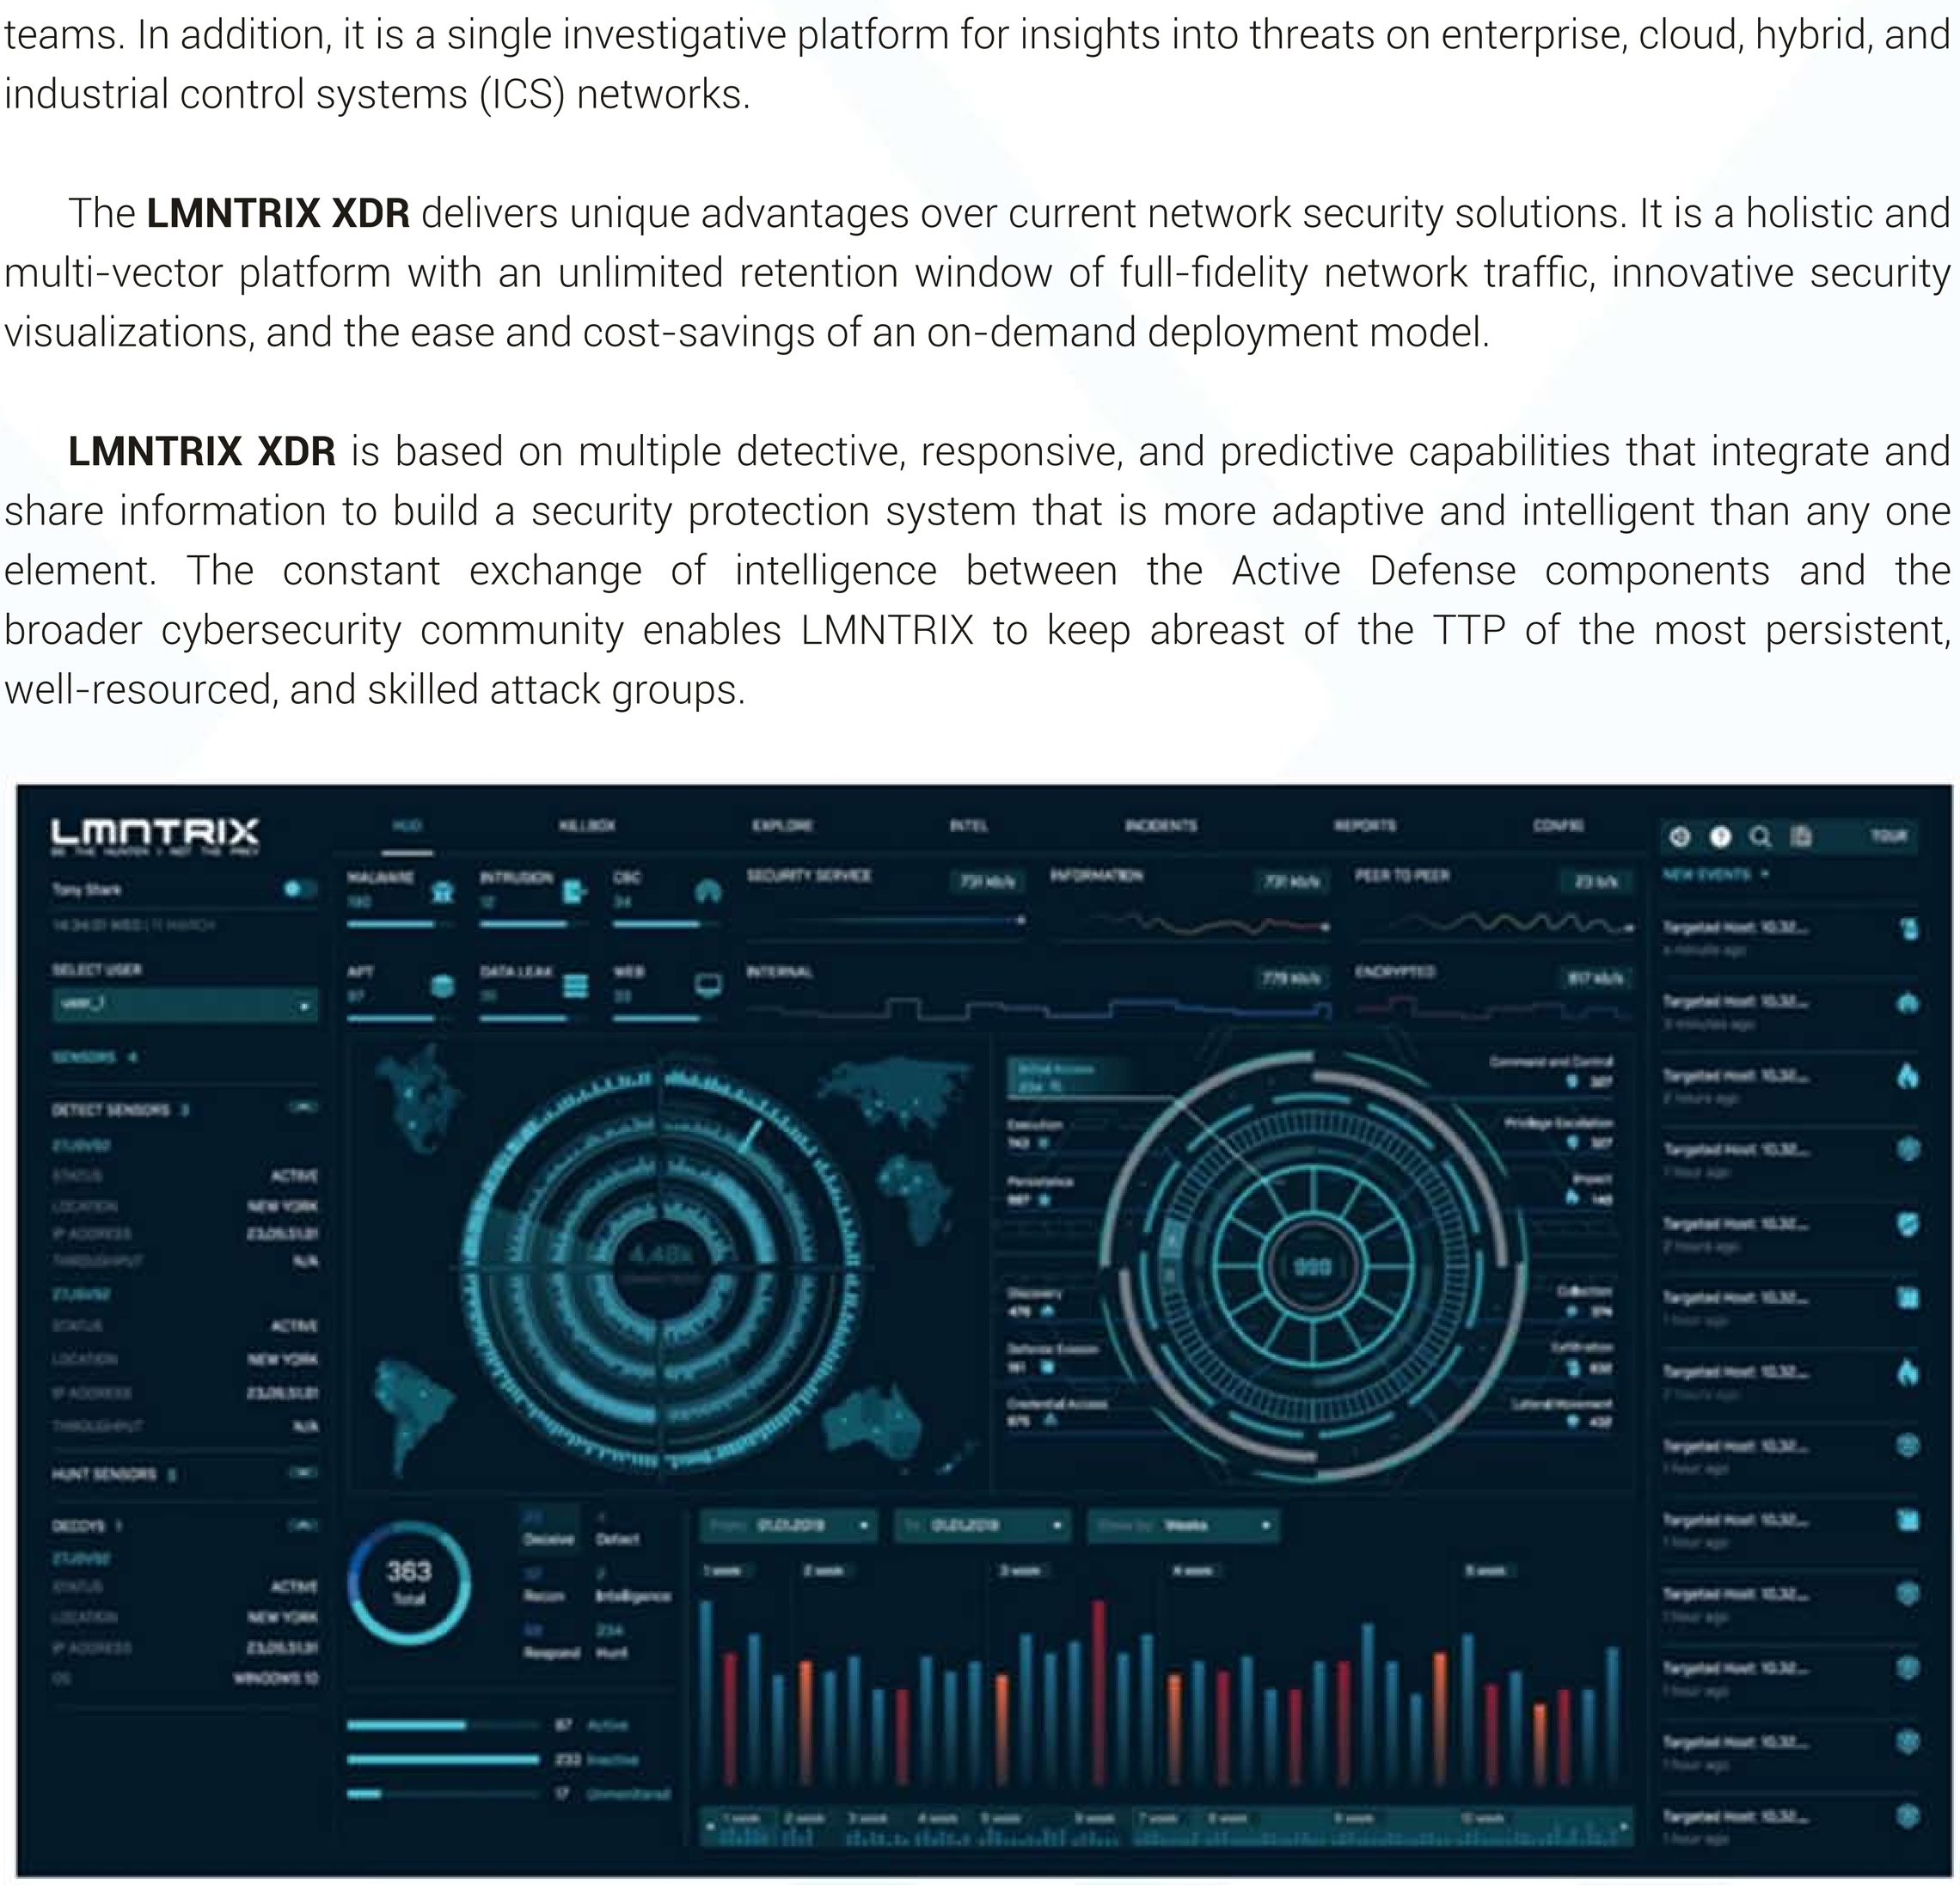

FIGURE 12 - LMNTRIX XDR

BE THE HUNTER | NOT THE PREY

LMNTRIX TECHNOLOGY STACK

The LMNTRIX Technology Stack is a powerful, proprietary threat detection stack embedded within the client environment behind existing controls. The technology stack comprises multiple detective systems, combining contextual threat intelligence and correlation, static-file analysis, user and entity behavior analytics (UEBA), and anomaly detection techniques to find threats in real time. In addition, it eliminates alert fatigue, determining which alerts to escalate through multi-platform consensus.

FIGURE 13 - LMNTRIX XDR Features

LMNTRIX CYBER DEFENSE CENTERS

LMNTRIX employs a global network of Cyber Defense Centers (CDC) comprising trained and certified hunters and intrusion analysts and provides constant vigilance and on-demand analysis of your digital assets and networks. Our intrusion analysts actively probe and monitor your networks and endpoints 24x7, using the latest intelligence and proprietary methodologies to look for signs of compromise. When a suspected breach is detected, the team performs an in-depth analysis of potentially affected systems to confirm the breach. Additionally, when data theft or lateral movement is imminent, our endpoint containment feature makes immediate action possible by quarantining affected hosts, whether on or off your corporate network. This significantly reduces or eliminates the consequences of a breach.

FIGURE 14 - LMNTRIX Cyber Defense Centre

LMNTRIX INTELLIGENCE

LMNTRIX INTELLIGENCE harnesses collated intelligence into a single platform to optimize collaboration and information sharing. Proprietary technology delivers earlier detection and identification of threats at every point along the attack lifecycle, making mitigating threats possible before material damage occurs.

- Threat actor communications interception identifies TTPs before they're used in anger and proof of concept attacks before full deployment.
- Correlation of over 850 million threat indicators against real-time network data deployed deep within customer networks.

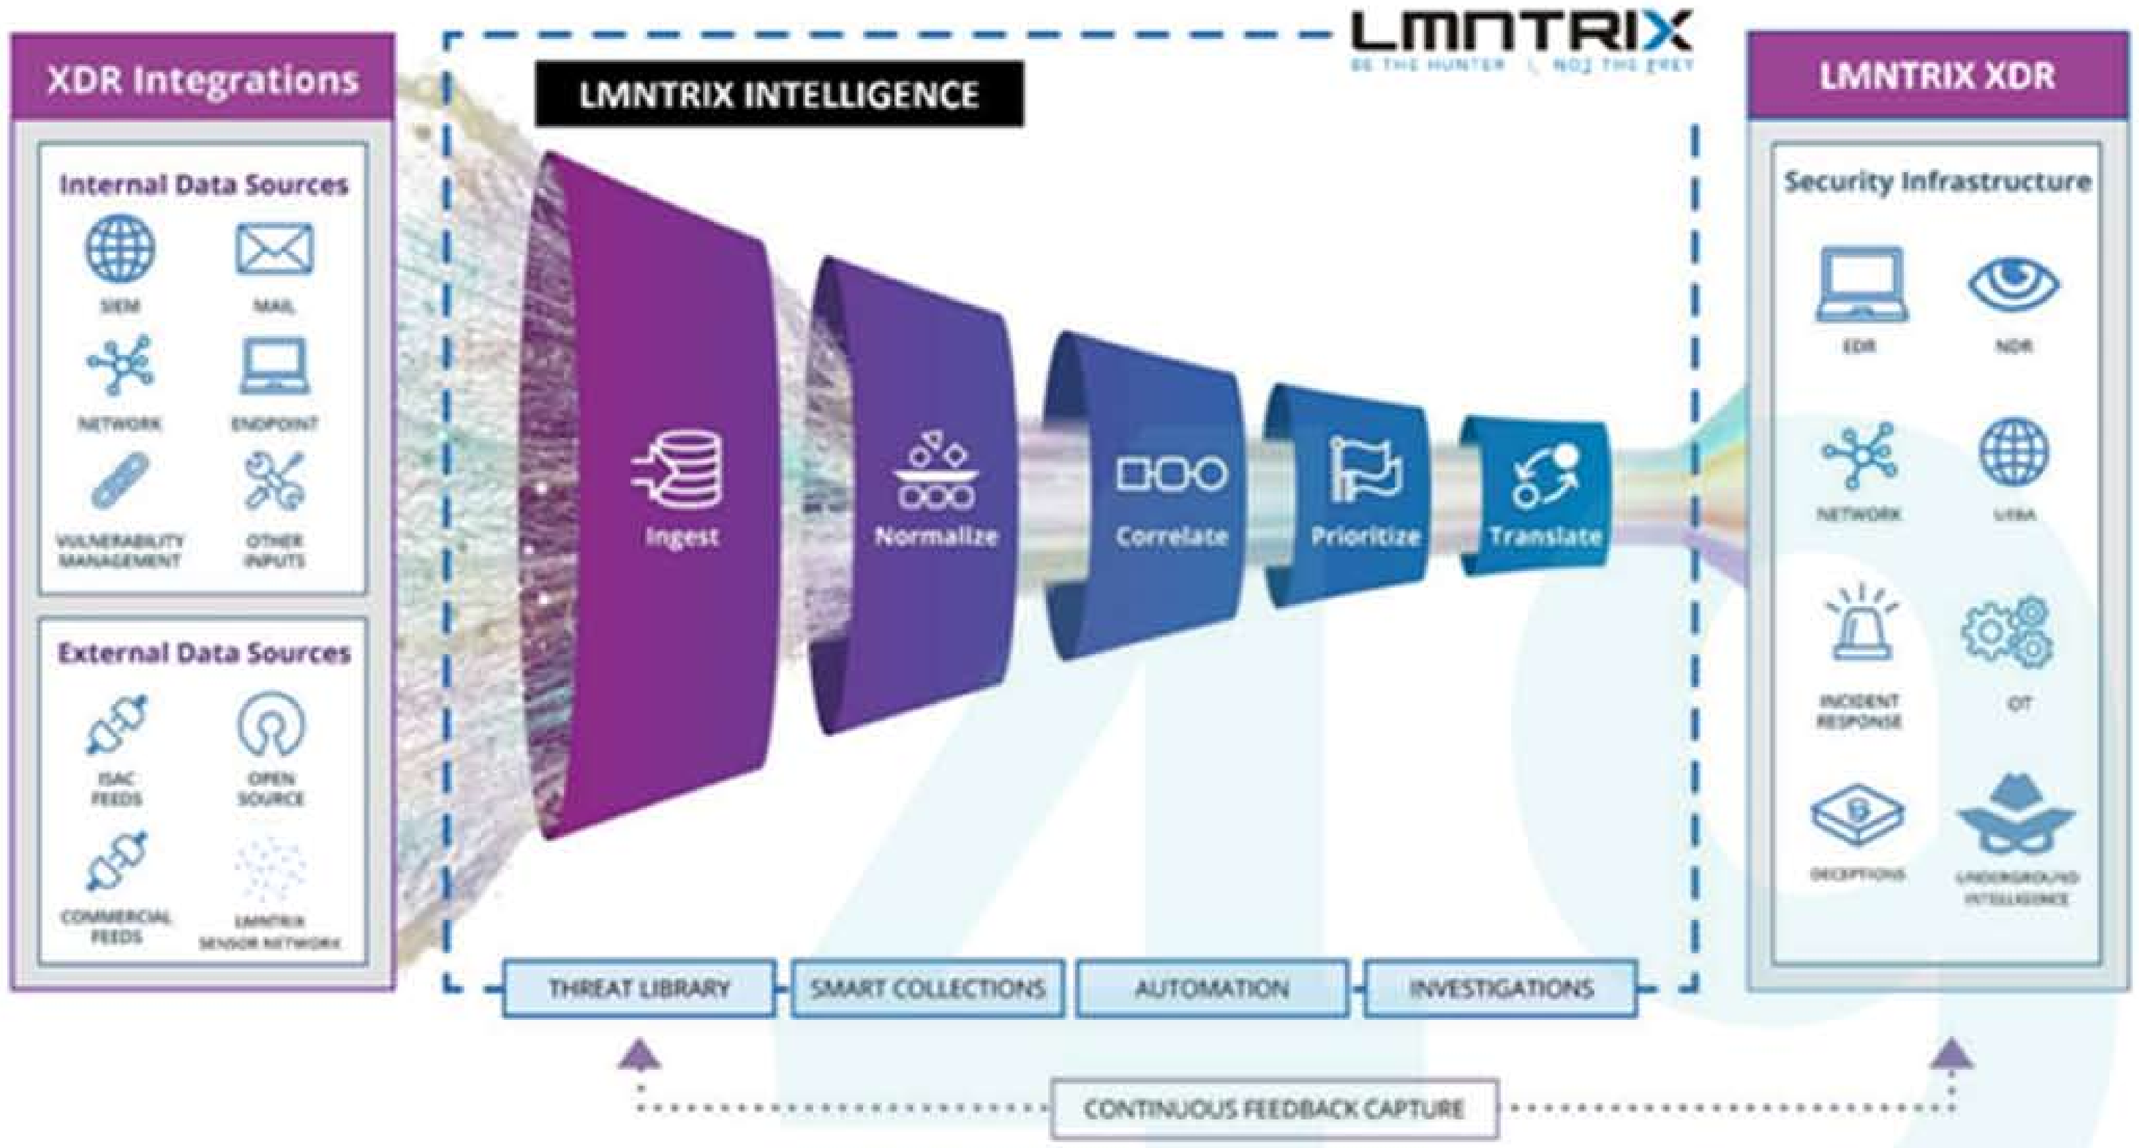

FIGURE 15 - LMNTRIX Intelligence Service

LMNTRIX

BE THE HUNTER | NOT THE PREY

LMNTRIX RECON

LMNTRIX RECON uses proprietary technology to detect indicators of attacks outside of a customer's systems by scanning for intelligence specifically relevant to the Client and evidence of stolen data.

- Monitors online black markets in the deep and dark webs for evidence of attack planning or the presence of illegally obtained data.
- Detects data leakage from malicious external and internal attacks and accidental disclosures.
- Monitors for reputational attacks and brand security alerting.
- Early detection of potential phishing domains allows preventative response.
- Aggregates unique cyber intelligence from multiple sources for analysis of cyber threats.
- Supports threat prioritization and remediation processes.

FIGURE 16 - LMNTRIX Intelligence Service

LMNTRIX

BE THE HUNTER | NOT THE PREY

ABOUT LMNTRIX

LMNTRIX has reimagined cybersecurity, once again turning the tables in favor of the defenders. We have cut out the bloat of SIEM, log analysis, false positives, and associated alert fatigue and created new methods for confounding even the most advanced attackers. We combine deep expertise with cutting-edge technology, leading intelligence, and advanced analytics to detect and investigate threats with great speed, accuracy, and focus. We believe that in a time of continuous compromise, you need a continuous response – not incident response.

As a company, we stand in defiance of the unwanted human presence within corporate networks by attacking the root of the problem—the adversary's ability to gain entry and remain undetected. Our real-time hunt operations identify signs of planned and active attacks and take action to neutralize them, forming the basis of our comprehensive Active Defense approach to limiting security exposure.

We are a partner who becomes an extension of your internal team, can augment your MSSP, or be a full-service SOC as a service security solution.

TO LEARN MORE
ABOUT **LMNTRIX** VISIT

https://lmntrix.com/

LMNTRIX USA.
333 City Blvd West,
18th Floor, Suite 1805
Orange, CA 92868
+1.888.958.4555

LMNTRIX UK.
200 Brook Drive, Green Park,
Reading, RG2 6UB
+44.808.164.9442

LMNTRIX INGAPORE.
60 KAKI BUKIT PLACE#05-19
EUNOS TECHPARK
+65 3159 0639

LMNTRIX HONG KONG.
14F, Manning House, 38-48
Queen's Road Central, Central,
Hong Kong
+852.580.885.33

LMNTRIX AUSTRALIA.
Level 32, 101 Miller Street,
North Sydney NSW 2060
+61 288.805.198

LMNTRIX INDIA.
VR Bengaluru, Level 5,
ITPL Main Rd,
Devasandra Industrial Estate,
Bengaluru, Karnataka 560048,
Email: sales@lmntrix.com
+91-22-49712788