AT&T Cybersecurity

A report from AT&T Alien Labs,
the threat intelligence unit of AT&T Cybersecurity

Published 13 May 2020

Data Analytics and Cyber Threat Intelligence

How Automation, Analytics, and Machine Learning Improve and Accelerate Threat Analysis

AT&T Business

Contents

Overview

The cybersecurity industry is increasingly producing enormous amounts of raw threat data. The sheer volume of information threat researchers must sift through makes it difficult to collect, analyze, and research that data in a timely manner. This in turn limits their ability to understand what data is valid and useful and whether threat artifacts will result in legitimate threat indicators.

In fact, it has been estimated that it would take 8,774 analysts working full time for a year to process the same amount of security event data that machine analytics can process in that same time frame.

Even as new threat intelligence tools and services emerge, relatively few enterprises are able to use those tools effectively due to the way threat intelligence and technology evolve. Threat actors are continually changing their methods of attack, and so the threat intelligence that supports detection must take new forms all the time to remain up-to-date.

In addition, cloud technology, 5G, edge computing, and the explosion of IoT devices is fundamentally changing the nature of threats and how defenders protect enterprises against them. Threat intelligence researchers are clearly facing a big data problem.

This paper considers why collecting and analyzing raw threat data today requires advanced analytics and machine learning (ML), in addition to human intelligence, to efficiently and accurately evaluate and interpret the volume of data that analysts must sift through on a daily basis.

It will also consider the stages of threat analysis that can be used to quickly turn raw threat data into curated threat intelligence that is fed into a variety of security technologies where it can be operationalized, such as a threat detection and response platforms. High quality, global threat intelligence is among the most powerful tools an organization has to defend against adversaries.

Introduction to cyber threat intelligence

AT&T Alien Labs™ defines cyber threat intelligence as the actionable information needed to continuously detect threats and prioritize response. This includes the ongoing collection, normalization, research and analysis, and correlation of threat data to drive the appropriate and most effective response. Threat intelligence includes more than atomic indicators (the tools threat actors are using, such as malicious IP addresses, URLs, or hash values). Threat intelligence also provides insight into the overarching behaviors of adversaries, including their motivations, intent, and techniques.

All of this information can be used to develop comprehensive attacker profiles that help researchers draw inferences to better predict future attacks and support resiliency in threat detection. By considering the overall tactics, techniques, and procedures (TTPs) of threat actors, and not just their tools, security professionals can use threat intelligence to its most effective and primary purpose: to drive resiliency against threats and ultimately protect the business, its data, and its customers.

The Lockheed Martin Cyber Kill Chain® model for attack analysis accepts threat indicators as the fundamental building blocks of intelligence. This includes any piece of information that objectively describes an intrusion. Threat indicators are threat data, pulled from many different internal and external sources, which have been validated as malicious or known to be malicious. They can be as simple as knowing that, for example, a particular bad actor prefers to target Windows machines. Or, threat indicators can be compiled to create attacker profiles that are as complex as knowing the various targets, aliases, and methods used by a highly successful hacking group such as Winnti, which is believed to have activity dating back to 2011.

Winnti's behavioral profile includes many variations of TTPs used in attacks that target multiple industries. For example, Winnti may use a phishing email to lure an IT employee into taking an action that ultimately results in their system being infected with malware. The malware, among other things, gives the adversary elevated access credentials and free reign to the business' network with a trusted VPN. Winnti can then move laterally using common network admin tools and can exfiltrate data through the business' trusted email services. These behaviors are just a few of dozens associated with Winnti. Researchers have developed a catalogue of attacks performed by this adversary group (or groups), including the common tools and techniques they use and relationships between attacks.

Because the threat landscape is always evolving, researchers and analysts must consider which technologies and methods are the most effective for analyzing, identifying, and containing threats in a particular moment.

Over the years, discussions on the most appropriate types of threat intelligence to use in detection and response have evolved. Some have declared the death of atomic threat indicators (such as IP address, file hashes, and domains) as detection tools, instead turning to behavioral-based approaches that identify and categorize the patterns and behaviors of malware and adversaries.

However, the increasing use of open–source tools among defenders has complicated malware attribution and clustering due to the fact that adversaries are using these same open–source tools to understand and adjust their attack methods. In addition, the emergence of commercialized cybercrime and crime syndicates has significantly impacted the level at which threat intelligence must be delivered due to malware families being modularized and sold on the black market as individual components that can be easily purchased and quickly used in an attack.

Threat researchers, therefore, must use multiple layers of intelligence to identify adversaries whose methods and behaviors will likely fluctuate or malware that may have many variations. These layers span the spectrum of simple indicators of compromise (IOCs) to more complex identification of common adversary TTPs and malware characteristics. By using layered threat intelligence, security professionals are able to better ensure resiliency in threat detection. (See figure 1.)

Figure 1: Using threat intelligence to identify patterns of behavior.

Atomic Threat Indicators
IP addresses, hostnames, file hashes, mutex values, and other attacker artifacts that historically have been used for intrusion detection and prevention systems.

Malware-Centric
Considers the relationship between files to create clusters with similar characteristics and behaviors.

Adversary-Centric
Considers adversary behaviors, such as intent, motivation, typical infrastructure, location, and patterns.

When it comes to identifying atomic threat indicators, research teams can use various forms of analysis to perform a variety of actions that would otherwise require manual work by a researcher. These tasks may include the daily extraction of threat indicators from dozens of vendor or government reports, alerts, articles/blogs, and social media.

Analytics and automation can also be used to cross-reference public databases for known, malicious URLs and IP addresses, scan web sites to understand the linkages between domains, update vendor signatures for new malware families, or scan files with multiple virus tools.

Some examples of threat indicators that can be automatically identified and extracted from reports, analysis, and unstructured data include:

- **CIDR Rules:** Classless Inter-Domain Routing, a set of IP standards that are used to create unique identifiers for networks and individual devices
- **CVE Number:** The Common Vulnerability Enumeration identifier of a vulnerability
- **Domains:** The domain name for a website or server
- **Email:** An email description, content, or headers
- **File hashes:** Strings of numbers and letters assigned to electronic data by a computer algorithm that provide a unique "digital fingerprint" of a file (e.g. MD5, SHA1, SHA256, PEHASH, and IMPHASH)
- **File paths:** The file system paths of known files and devices (i.e. the complete location or name of where a computer, file, device, or web page is located)
- **Hostnames:** The subdomains for a website or server
- **MUTEX name:** A mutual exclusion object (a program object that allows multiple program threads to share the same resource, but not simultaneously)
- **IP addresses:** An IPv4 or IPv6 address that identifies each machine/device using the Internet Protocol (IP) to communicate over a network

- **URI:** The Uniform Resource Indicator (URI) describing the path to a file hosted online
- **URL:** The Uniform Resource Location (URL) summarizing the online location of a file or resource

Simple threat indicators are a useful starting place as a first line of defense and in building malware and threat actor profiles. However, they should not be relied on alone. These are the tools that threat actors can (and do) change frequently and quite easily, often using automation themselves.

The Pyramid of Pain (see figure 2), first published in 2013 by David Bianco, articulates that while a wide variety of indicators should be used to detect threats, not all threat indicators are created equal. The atomic indicators at the bottom of the pyramid (such as IP addresses and hash values) are easy to acquire and relatively easy to feed into security platforms. However, these are also very easy for threat actors to change.

For more resilient threat detection, it is important to go beyond atomic indicators to also identify higher level indicators, such as network and host artifacts. This can include anything that indicates an attacker has been in a network or on a host using either signatures-based IDS (which looks for specific and known patterns such as byte sequences in network traffic) or anomaly-based IDS (which monitors for abnormal system activity to detect unknown or emerging attacks, often via machine learning). Additionally, threat analysts can make use of Yara rules, which describe the patterns of a particular strain or family of malware or shared code. These more complex IOCs are more difficult to acquire, but they are also more difficult for attackers to change.

The most resilient threat detection comes when both simple and more complex indicators are combined to define common adversary behaviors (or TTPs). Threat intelligence that delivers on this is able to describe how a threat actor goes about accomplishing their mission, from reconnaissance to data exfiltration and every step in between.

By focusing in on the behaviors and not just the tools adversaries are using, threat researchers are identifying the things that are the most difficult and expensive for attackers to change. This is why threat intelligence frameworks such as MITRE ATT&CK™ outline the common tactics and techniques used by attackers versus atomic indicators.

Figure 2: Pyramid of Pain illustrates the degree of difficulty for attackers to change.

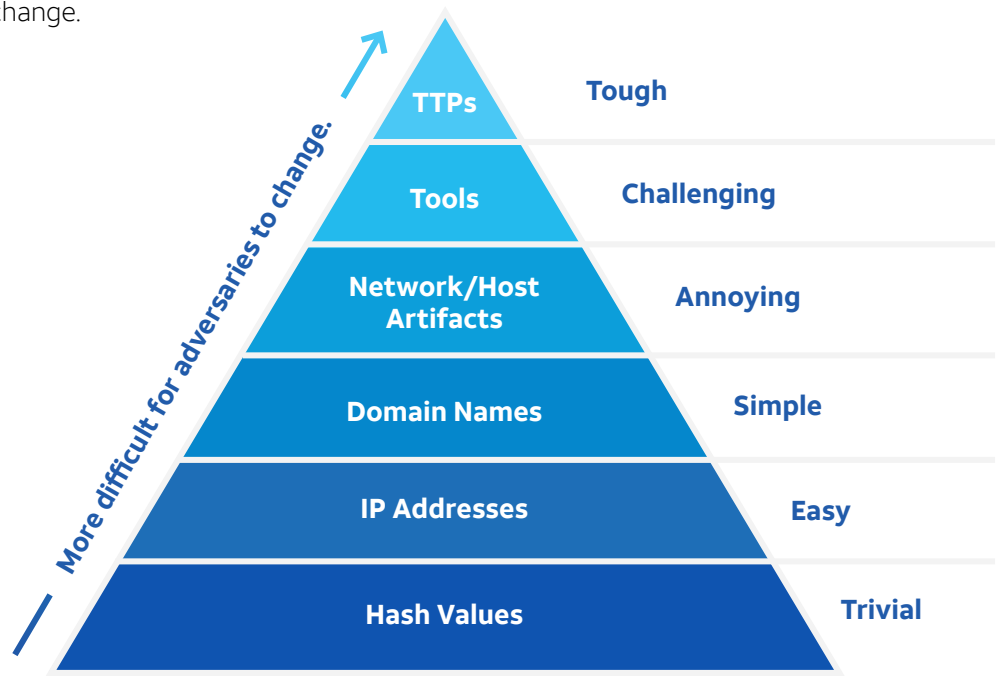

How data analytics and machine learning power threat analysis

Before diving into threat analysis, it is helpful to first clarify some terms. What do artificial intelligence and machine learning really mean, and how do they relate to each other? Simply put, AI brings decision-making capabilities to computers. AI is not new, but it has become more mainstream through the economics of cheap computing, including hardware, storage, and compute power.

Machine learning, a subset of artificial intelligence, is the ability of machines to automate a learning process and is used to identify patterns and make predictions. In addition, ML includes deep learning, which is used in highly automated systems where a critical mass of data is available for training such as image recognition, speech recognition, and more. Figure 3 shows how AI and ML relate to each other.

Figure 3: Artificial intelligence and machine learning.

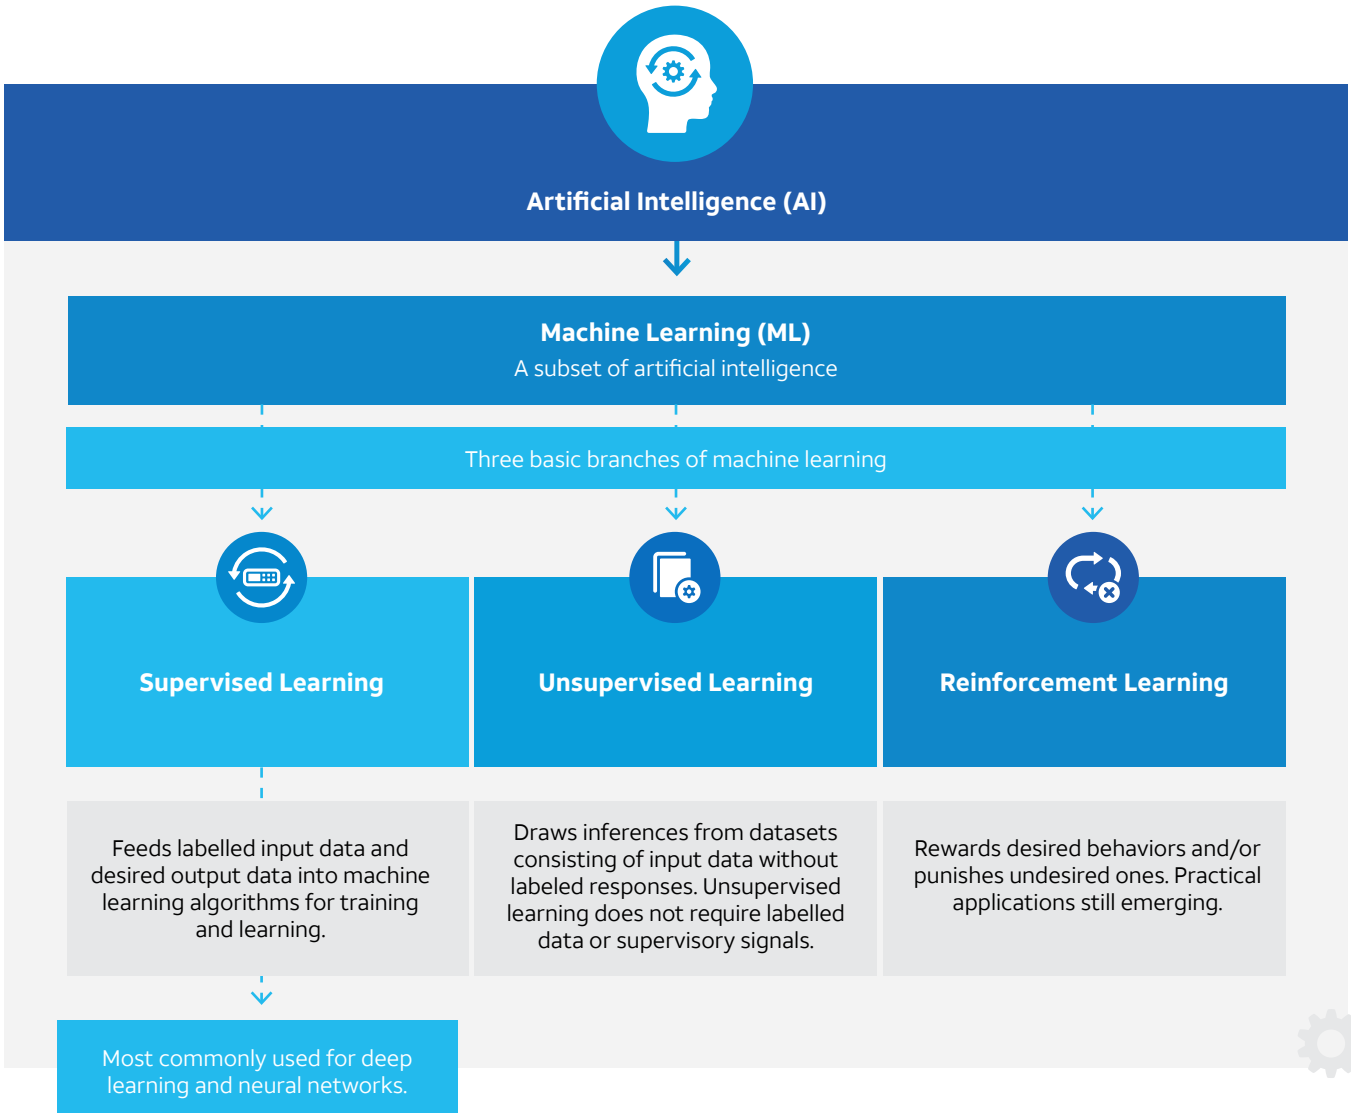

Machine learning models

Because threat actors re-use and evolve known malware, machine learning models are often used to create malware clusters that can detect and predict the behaviors of malware families. This helps to speed classification and identification. As noted, ML is a computer's ability to learn without being programmed. Ultimately, the objective is to move to a fully automated state, in which rules, thresholds, and metrics are fine tuned as the data changes.

The adoption of cloud technologies has been particularly useful for collecting data that can be fed into ML models as well as for feedback to retrain the data (such as reports of false positives and other misclassification). Additionally, the cloud enables other features in a supervised machine learning model, such as how many endpoints have seen a file.

By definition, all ML techniques use data to learn. It is often the quality of that data that can make a difference between the successful identification of a threat or a false positive.

Unsupervised machine learning is an approach in which data scientists use the dataset, a collection of unlabeled examples, to train algorithms to find patterns and make predictions about new data. Unsupervised ML is used to "draw inferences from datasets consisting of input data without labeled responses."

The process creates a set of classes that the model "thinks" are relevant and creates a baseline for normal behavior. The model then finds patterns that deviate from the norm. Models are not trained ahead of time, and because of that, false positives can be frequent and results often need to be validated. This takes time, and when time is of the essence, it can make more sense to rely on more traditional IOCs such as signature- or anomaly-based IDS. For this reason and because unsupervised ML is still maturing, it is not as widely used as supervised machine learning.

Still, unsupervised machine learning has its place. For example, it can be used to cluster malware families (typically in concert with supervised ML models). That is, unsupervised machine learning helps discover the inherent groupings or associations in the malware data.

As another example, AT&T Alien Labs uses unsupervised machine learning to speed IOC extraction from threat data submitted to the Open Threat Exchange (OTX). OTX is a crowd-sourced platform where users create "pulses" that contain information about a recent cybersecurity threat. A pulse consists of threat indicators and links to blog posts, white papers, reports, and other files with attack details. It typically contains a link to the full content (a blog post, for example) plus key metadata that a user or researcher can manually extract from the content (the malware family, target of the attack, etc.). Using unsupervised machine learning, the Alien Labs team can automate this extraction process and enrich the pulse with additional information identified by the Alien Labs systems. This increases the efficiency and speed of threat intelligence collection.

THE OPEN THREAT EXCHANGE (OTX)

AT&T Alien Labs Open Threat Exchange™ (OTX) is a free, open-source and global community of more than 140,000 threat researchers and security professionals in 140 countries who actively research and share up-to-date threat intelligence on indicators of compromise (IOCs) as well as the TTPs that threat actors use to orchestrate attacks. The community is free to join and provides valuable tools such as the ability to download IOCs via an API, free threat analysis, and auto-extraction of IOCs from dozens of files including PDFs, emails, and more.

Supervised machine learning focuses on classification and prediction, based on known properties previously learned from the training data. According to Gartner, "supervised learning is the most popular and most frequently used type of machine learning in enterprises, because it has been proven to work well in many business scenarios. It works by feeding input data and desired output data into machine learning algorithms."

Most next-generation anti-virus (NGAV) and other security tools such as UEBA also use supervised machine learning. And, because the industry already has access to a large volume of labeled malware files, supervised models typically have enough training data to make highly accurate predictions and classifications with less need for validation and tuning of the data. For example, data may be classified as malicious, suspicious, benign, or unknown.

Supervised machine learning can also be used for detecting things like common line obfuscation, malicious power shell obfuscation, and anomalies in time series. It can also predict domain-generation algorithms or classify domains to remove false positives in automated command and control (C&C) extraction.

All the while, the model can learn from the data it is ingesting, creating a **neural network**. Neural networks are becoming more frequently used in supervised learning problems. They are able to learn more complex patterns and can improve results over other algorithms because they can train larger models with more data.

From threat artifact to threat intelligence

Threat research labs typically use a combination of analytics and machine learning to process the volumes of threat data they ingest. (Some advanced research labs are ingesting 20 million threat artifacts* per day or more.) Figure 4 (on the following page) shows an example of how threat data may go through various stages of analysis.

With such a huge volume of data, researchers must use these tools to automate analysis and help narrow their field of research, so they can then focus in on the threats that need further investigation. Machine learning also helps identify the patterns within indicators that could reveal new or evolving malware or adversary TTPs.

Threat identification and analysis can be broadly categorized as static analysis, dynamic analysis, and hybrid analysis.

Static analysis is the process of analyzing malware or binaries without actually running the code. It can be as simple as looking at metadata from a file and can range from disassembling or decompiling malware code to analyzing the intermediate representation of program source code. Static analysis is done through a variety of techniques, including *signature-based* or *heuristic-based* techniques.

For example, using a *signature-based* detection technique, the malware analyzer is looking for known pattern matching in the signatures (the bit of sequence injected in the application program by the malware writers that uniquely identifies a particular piece of malware).

Heuristic detection takes this one step further. In this technique, instead of looking for a particular, known signature, the malware detector is searching for patterns that might indicate a certain behavior. Because heuristic detection is not based on a specific signature being known at a single point in time, it becomes easier to detect new variants of malware that have not yet been identified.

Heuristic techniques may include looking for commands to delete or harm other files, or for variants of known, malicious signatures. Other

Figure 4: An example of systems that can be used to analyze threat data.

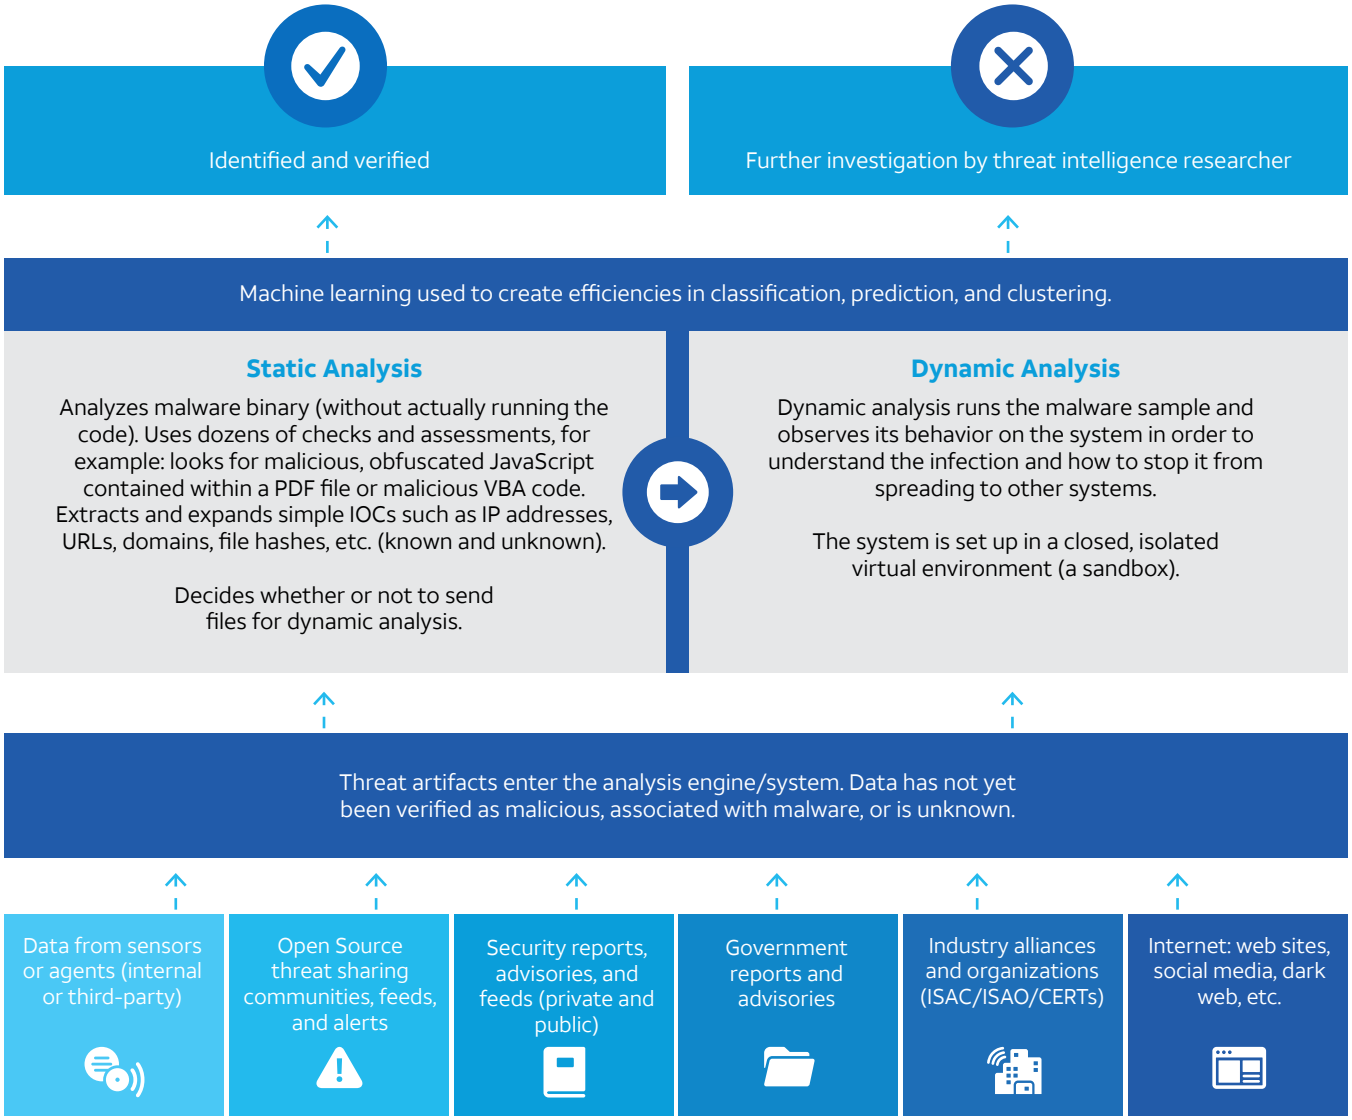

examples include looking for patterns that might indicate obfuscated JavaScript contained within a PDF file or malicious VBA code. (Visual Basic for Applications, VBA, is the programming language of Excel and other Office programs.) Based on the results, static analysis agents can determine whether or not to send files for dynamic analysis.

Dynamic analysis can be use to further analyze and identify malware samples not seen before (or variants not recognized) by running a sample and observing its behavior. Looking at the behavior of the malware and its side effects helps to understand the infection and how to stop it from spreading. A dynamic analysis system is set up in a closed, isolated environment — a virtual machine or "sandbox." In this environment various tools are used, such as process monitor and Sysmon (System Monitor), to see what kinds of artifacts the malware produces when it is run. Machine learning can be used to perform multiple tasks during dynamic analysis, such as taking a domain from the sandbox and writing signature variants for it. It is very easy for attackers to switch to new domains, and ML gives defenders a leg up by helping to predict the possible variations on how they might do so.

The ML model can be trained with information about malicious domains as well as benign infrastructure. This is important, because attackers could be using legitimate infrastructure such as Gmail™ for communications. If there are samples in the sandbox communicating with domains, the model can tell whether that is a malicious domain or benign infrastructure, and the model is able to filter the results accordingly before they can be used as indicators of compromise.

A hybrid approach combines both static and dynamic analysis, first checking for a known malware signature and if it is present in the code, then monitoring the behavior of the code in a sandbox. As an example, for malware that creates a registry entry, the machine learning model is trained to look at static data to predict whether a specific file has a known behavior (label). Only unknown files are sent to dynamic analysis. In this way, hybrid analysis enables more efficient processing.

Extracting and expanding threat indicators

As an example, static analysis is used to speed up the process of extracting threat indicators from files, web pages, etc. It is also used to augment the information about those indicators with important contextual details, such as country of origin, targeted industry, and malware name.

It can also be used to perform scans with multiple anti-virus tools and perform dozens of automatic checks to look for malicious code or associations with known, malicious URLs or IP addresses, as well as other threat indicators. These static checks are typically performed against a database that an organization's research team builds or acquires. Suspicious samples are routed to other resources in the analysis engine for further review.

To illustrate this, files submitted to the AT&T Alien Labs Open Threat Exchange (OTX), are automatically run through a malware and threat analysis engine, which includes multiple layers of automated checks, data analytics, and machine learning. Files and URLs are quickly analyzed, first with static analysis and then, depending on the file type, will go to a sandbox for dynamic analysis, including an assessment of network activity. For URL submissions, the analysis engine scans the URL looking for threat artifacts such as a suspicious
file, and if found, those artifacts are submitted for further analysis.

The tools used in static analysis help identify and enrich the information that is then used to build comprehensive profiles of malware and adversary TTPs. These tools vary by research team, but some examples of actions performed include:

- Checking public databases for bad URLs and IP addresses
- Acting as a honey client to identify whether a domain is distributing malware
- Checking the Whois records for ownership of domain names
- Scanning web sites to see if domains link to each other
- Identifying malicious IP addresses
- Identifying if a URL tries to download malware
- Notifying about a new signature for a new malware family
- Converting malware into a higher-level programming language to help unmask the inner workings and functions of its code
- Storing exploit details (for example, whether an exploit is being actively used in the wild)
- Monitoring social media, such as Twitter®, to detect for alerts to threats

Malware classification and clustering

In addition to static analysis, and depending on the information collected from that stage, the analysis engine may direct suspect data to dynamic analysis and ML systems for further assessment. Machine learning techniques use trained data models of previously identified malware types, which can be used to identify if a particular threat is similar to others within a malware cluster (malware clusters

are comprised of malware families with similar behaviors). For example, is the same file modified or injected with a similar process? Machine learning analyzes the results to detect and classify malware families and behaviors, turning what might have been a million threat indicators into 20-30 clusters that can then be further analyzed by a threat research team.

Clustering malware into sets also enables the analysis engine to more quickly identify additional malware samples found in the wild, including whether a sample is a completely new instance or a variant of a well-known family. With this information, new detections can be quickly created for new variations targeting the environment. It also makes it easier to derive generalized signatures, implement removal procedures, and create new mitigation strategies that work for a whole family of malware. Malware clusters narrow the field of information that threat researchers need to examine and speed their ability to validate, evaluate, and interpret that information so they can then turn the data into rich, actionable threat intelligence that can be quickly operationalized.

Summary and next steps

With the availability of data analytics, automation, and machine learning models, threat researchers have the tools needed to help deal with the enormous (and growing) volume of threat data. In addition, these tools give researchers an efficient means for curating threat indicators and quickly turning that information into valuable and operational threat intelligence.

However, it is an expensive and resource–intensive proposition for a company to create its own analytical systems and ML models for cyber threat intelligence. To build an effective and automated analysis system for threat intelligence requires experience, a sizable monetary investment, dozens of hard-to-find researchers and data scientists, and access to an ongoing feed of global threat data.

This is why many companies turn to third-parties for their threat intelligence. When looking for external assistance, companies should consider a vendor with:

- Access to a large, globally diverse pool of threat data
- Strong understanding of the effectiveness of various processes and technologies for threat analysis
- An experienced research team staffed with threat intelligence researchers, threat analysts, and data scientists
- Proven history of identifying, analyzing, detecting, and containing new and evolving threats
- Multi-layered threat intelligence that includes the full spectrum of IOCs and adversary TTPs

Additional resources

AT&T Alien Labs Website

AT&T Alien Labs Threat Intelligence Datasheet

Malware Analysis by Open Threat Exchange

AT&T Cybersecurity

AT&T Cybersecurity helps to reduce the complexity and cost of fighting cybercrime. Together, the power of the AT&T network, our SaaS-based solutions with advanced technologies including virtualization and actionable threat intelligence from AT&T Alien Labs and the Open Threat Exchange™, and our relationship with more than 40 best-of-breed vendors, all accelerate your response to cybersecurity threats. Our experienced consultants and SOC analysts help manage your network transformation to reduce cybersecurity risk and overcome the skills gap. Our mission is to be your trusted advisor on your journey to cybersecurity resiliency, making it safer for your business to innovate.

www.cybersecurity.att.com

About AT&T Alien Labs

AT&T Alien Labs, the threat intelligence unit of AT&T Cybersecurity, delivers continuously updated threat intelligence to the cybersecurity products and services our customers trust to protect their business. Alien Labs includes a global team of threat researchers and data scientists who, combined with proprietary technology in analytics and machine learning, analyze one of the largest collections of threat data in the world. Alien Labs goes beyond delivering threat indicators to performing deep, qualitative research that provides insight into adversary tactics, techniques and procedures (TTPs). By identifying and understanding the behaviors of adversaries (and not just their tools), we can help power resilient threat detection, even as attackers change their approach or an organization's IT systems evolve. Our threat intelligence is integrated directly in to the AT&T USM™ platform in the form of correlation rule sets that are updated daily for on-premise and multi-cloud environments, including endpoints. This direct integration helps our customers shorten the time from public disclosure of a threat to response and containment.

References

Alaybeyi, S., Linden, A., den Hamer P. (2020). Types of Machine Learning for the Enterprise. Gartner.

Bianco, D. (2014). The Pyramid of Pain. Retrieved from http://detect-respond.blogspot.com/2013/03/the-pyramid-of-pain.html

The Cyber Kill Chain. Lockheed Martin. Retrieved from https://www.lockheedmartin.com/en-us/capabilities/cyber/cyber-kill-chain.html

Gopalakrishnan, S. (2017). Data Science & Machine Learning in Cybersecurity. AT&T.

Haniyur, S. Entity Extraction for Threat Intelligence Collection. Retrieved from https://cybersecurity.att.com/blogs/labs-research/entity-extraction-for-threat-intelligence-collection

The MITRE ATT&CK® Framework. Retrieved from https://attack.mitre.org

Robertson, C. (2013). Indicators of Compromise in Memory Forensics. SANS Institute.

Zelonis, J. (2017). Achieve Early Success In Threat Intelligence With The Right Collection Strategy. Forrester Research.

What is Threat Intelligence? Recorded Future. Retrieved from https://www.recordedfuture.com/threat-intelligence/

* Based on threat artifacts contributed to the Alien Labs analysis systems on a daily basis.

AT&T Business