Business white paper

Collaborative Defense Enriched by Dynamic Analysis

Table of contents

Our adversaries are collaborating

Today's dynamic threat landscape requires that enterprises be more vigilant than ever when it comes to protecting sensitive data and information resources. The evolving nature of threats is a top security intelligence challenge that organizations face, and targeted attacks are on the rise. Cyber attacks are increasingly more sophisticated and organized. Adversaries specialize in different aspects of an attack and collaborate to achieve their objectives. In light of these advancements, current security research, detection and response capabilities are inadequate and regular headlines about successful attacks demonstrate the impact on businesses.

A notable example is the $45 million USD ATM heist that spanned 2012-2013. USA Today reported on May 10, 2013 that an international gang of cyber thieves stole $45 million USD from thousands of ATMs with a carefully coordinated operation conducted in a matter of hours. There were two separate attacks on two different Indian credit card companies and two different Middle Eastern banks. In December, attackers stole $5 million USD and in February, they snatched $40 million USD in 10 hours using thousands of transactions. If the first bank had informed their peers, perhaps the second, larger theft could have been prevented.

To combat collaborative attackers, enterprises must also collaborate to create a united threat intelligence defense. To understand how, let's first look at the stages of a cyber attack.

The five stages of a cyber attack

There are distinct phases of a cyber attack. It's important to understand that these stages may not always happen in a linear flow—some stages may occur at the same time or even repeat multiple times. The attack life cycle represents how far along an adversary has progressed in the attack and gauges the damage that has occurred.

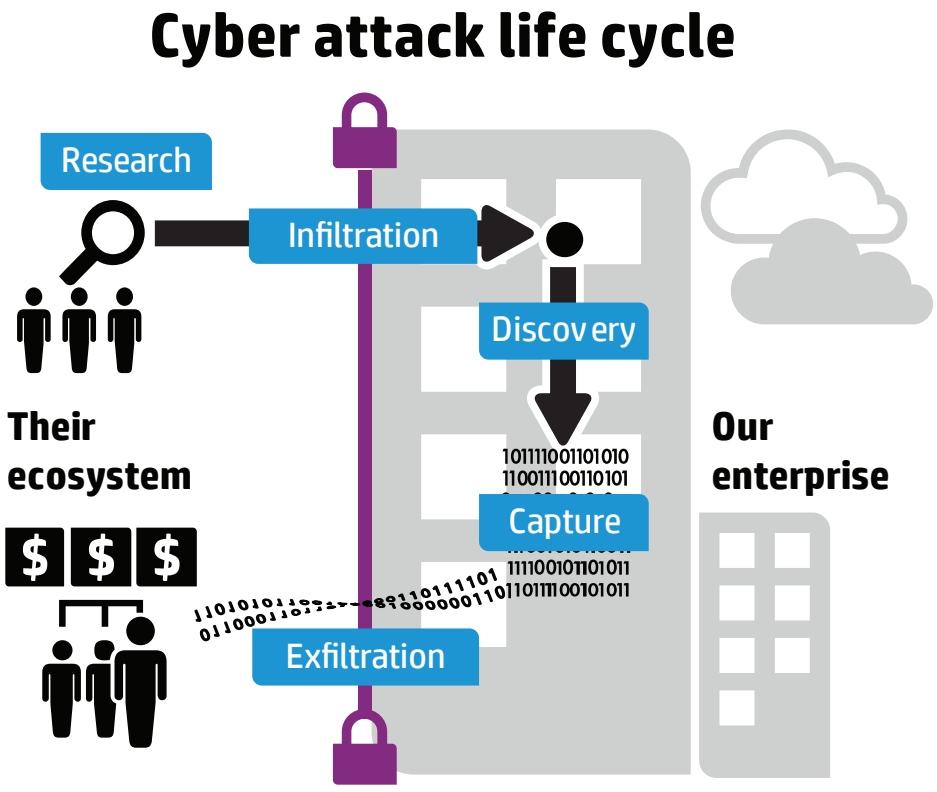

1. Research: Research drives the identification and selection of targets. The goal is to gather the necessary information to penetrate network and perimeter defenses in an undetected manner.
2. Infiltration: Infiltration is the act of placing a malicious payload into a delivery vehicle and delivering the weaponized bundle to the target environment. Adversaries often target vulnerable software systems at this phase and will attempt to clear any evidence of malicious activity.
3. Discovery: Once inside the enterprise, adversaries probe to identify intellectual property, sensitive data, or other targeted assets. When found, the adversary plans the necessary steps to take control of, or capture, the assets.
4. Capture: Until this phase, adversaries may have gained access to a broad range of systems and identified assets to target for exfiltration. In the capture phase, adversaries capture the assets and prepare the asset for exfiltration.
5. Exfiltration: The exfiltration phase involves shuttling the captured assets outside of the victim organization and typically marks the successful completion of a breach.

There is no longer a single enemy

Today, many actors specialize in narrow abilities related to different phases of the attack. Threat actors planning large operations combine the abilities of multiple specialized groups to achieve their goals, which effectively creates an ecosystem governed by market forces. This ecosystem is efficient at creating, sharing, and acting upon its security intelligence-utilizing sophisticated exploits and malware in long-term, multi-stage attack campaigns that target specific organizations. By combining a multitude of technical capabilities from different specialists, adversaries lead advanced attacks that are difficult, if not nearly impossible to detect. These adversaries pose more than a singular threat-they have developed into a marketplace.

A bigger picture

The marketplace adversary demonstrates persistence with a larger strategic goal than any individual. The marketplace adversary is not deterred by early failures and there is a high probability that they will attack the same target repeatedly, utilizing diverse techniques, until they succeed. If the adversary is discovered, they will regroup and try again.

A security research study conducted by HP Labs and HP TippingPoint of HP Labs in Princeton, NJ, together with Marc Eisenbarth of HP TippingPoint, analyzed data from over 35 billion alerts recorded over a 5 year period by HP TippingPoint Intrusion Prevention System (IPS) devices located in over 1,000 customer networks worldwide. The collective data provides a rich view into the nature of attacks, both external and internal, across diverse networks. This cyber security research reveals that the majority of customers were initial targets of less than 10 attacks, while three customers were early targets of over 100 attacks. These repeated, persistent attacks, suggest that is useful to share information about early attack sightings with other organizations so they can determine their risk and any potential mitigations. (ACM Badgers 2012, Examining intrusion prevention system events from worldwide networks).

An adversary may attempt to use a common exploit to compromise a target system making it possible to connect past attacks conducted by the same adversary. For instance, different proxies might be used to launch an attack, but with the same weaponized exploit payload. These seldom-changing properties of attacks help define the adversary's behavioral profile. It's recognizing and detecting these techniques that expedites the discovery of additional attacks.

Protecting your data through collaboration

In order to counter attacks from marketplace adversaries, government agencies and private organizations must learn from the adversaries and collaborate by sharing security intelligence. Organizations must be able to respond quickly and effectively in order to beat the adversaries at their own game. For this to be feasible at scale, the industry needs a system that allows organizations to collaborate and share threat intelligence information in a secure, confidential, and timely manner.

Cyber security research and open collaboration can help organizations in a community avoid falling victim to attacks other community members have already experienced. In addition to details of an attack, a platform for sharing threat intelligence can be used to distribute the effort of analyzing evidence and sharing mitigations once they are developed. In some cases, better security intelligence can also be gleaned by combining data from various organizations that cannot be derived from any individual organization.

Moving from Collaboration to Analysis and Action

More and more stakeholders—security vendors, security analysts and researches, even governments—are understanding the value of threat intelligence exchange and collaboration. From a process standpoint, community-driven efforts such as STIX and TAXII, open standard protocols which enable automated information sharing, have already laid the groundwork for information sharing initiatives. Vendors, in turn, are falling in line with new announcements of collaborative threat intelligence systems.

However, at this phase of the threat information exchange lifecycle, simply sharing and consuming data should be viewed as “table stakes.” The true promise of data sharing lies in the ability of threat intelligence platforms to analyze the data and ultimately provide concrete guidance to the necessary action. A platform providing true actionable intelligence would:

1. Evaluate indicators come from multiple sources, each providing a unique view on the threat
2. Leverage collaboration and social interaction to link artifacts and layer additional context around indicators
3. Interact with an intelligent system to determine which threats are important to the analyst, providing the crucial aspect of relevance
4. Distill the analysis to trusted score, making the indicators actionable

Platforms of tomorrow will evolve from the leading edge state to automatically apply the resulting tactically.

Experience to build upon

Sector-wide sharing
Two examples of organizations that have been successful at facilitating defensive collaboration include the Financial Services Information Sharing and Analysis Center (FS-ISAC) and the Information Technology-Information Sharing and Analysis Center (IT-ISAC). There are several ISAC organizations, aligned by industries to help defend against threats. Members must be vetted to participate.

The ISAC organizations understand that when an attack occurs, an early warning can be the difference between business continuity and catastrophe. In the case of a threat, members of these groups are given information that is specifically designed to help protect critical systems and assets from physical and cyber security threats. In addition, these groups have specialized forums dedicated to risk management for participating firms. Members participate in collaboration efforts to strengthen the IT infrastructure through cyber information sharing and analysis. Collaboration helps their members improve incident response. These groups demonstrate some of the most recent examples of successful collaboration yet they still do not solve the problem. Much of the information exchange remains a manual effort.

Collaboration for vulnerability discovery

HP's Zero Day Initiative, or ZDI, is another example of information sharing with a goal of helping companies identify threats early in order to block malicious traffic that may take advantage of new vulnerabilities. Launched in 2005, ZDI's purpose is to augment HP DVLabs research with the additional zero day research from external researchers by offering financial rewards for submitting novel, verifiable vulnerabilities. While the ZDI team works with the impacted vendor to develop a patch, customers are protected by filters that, when possible, protect against the vulnerability. The manner in which the ZDI filters are developed minimizes the chance that adversaries will discover the details of the vulnerability before a patch is released while maximizing the protection provided to the community. While highly effective, ZDI filters only benefit HP TippingPoint customers.

Broader community-based collaboration and analysis

From a product standpoint, consider HP Threat Central, which is an open and automated cloud based platform for security intelligence, enabling customers to consume and share community driven intelligence. The platform combines the benefits of secure sharing along with the delivery of near real-time analyzed and actionable results.

A typical example of how such a system can optimize reactions for network defenders is the reconnaissance use case. Consider an attacker using the IP address 1.1.1.1 to probe a network for "open doors" through which to attack. That activity is picked up and logged by the target's intrusion protection system, but in a reconnaissance mission, an attacker very often uses a different IP address from which to perpetrate an actual exploit. Assume that IP address is 2.2.2.2. The activity recorded during reconnaissance is of no help preventing the actual attack, and because there is no sharing or analysis being applied to the observations as they occur, multiple companies can fall victim to the attack before one or both IP addresses are recorded as threats and shared via popular feeds. Very often, by the time the attack is properly recognized, it is no longer actionable intelligence because the attacker has moved on to new set of IP addresses from which to conduct his activities.

Now, consider the case where you have the network effect of a broad community of similar targets, all sharing threat intelligence through Threat Central. A first, even second target may see the behavior of 1.1.1.1 and record it as a low scored security event. But as the counts go up, and analysis within Threat Central indicates that there is a connection between the reconnaissance activity occurring from 1.1.1.1 and the attack coming from 2.2.2.2, the score on both indicators increases and a link is established between them. This link makes 2.2.2.2 a relevant indicator for anyone that sees 1.1.1.1 in similar context, meaning that this community of companies under attack receives the analyzed and actionable intelligence at close to real-time that they should block activity seen from 1.1.1.1 and 2.2.2.2. even if they have not yet observed either IP address on their networks.

The benefits of sharing security intelligence are widely recognized, but the practice has been ad hoc in terms of both the content shared (focused narrowly on a few security indicators) and the means by which the content is shared (generated and consumed manually via email and Web forums). Because of this, follow-on actions for the recipients of new intelligence are often unclear as are guidelines for what can be shared with whom and when. These concerns often make organizations unwilling to participate in information sharing. Threat intelligence feeds can be difficult to interpret and digest and the volume can be overwhelming. As security standards evolve and concerns grow about the instability of critical infrastructures, there is a strong push to develop a trusted framework for exchanging security intelligence.

Challenges to collaboration

Sharing security intelligence raises a number of important questions—for instance: What type of data will be shared and what will remain outside the scope of the exchange? Where and how will information be generated and consumed by participating organizations? How will the data be represented, stored, searched, and analyzed? How trustworthy are the results?

Example: What to share and how?

Consider the case of a spear phishing email. In a classic spear-phishing attack, the victim might receive a seemingly legitimate email that includes a malicious attachment or that directs the victim to a malicious webpage in a guise to learn logon credentials or to utilize a browser exploit to download malware to the victim's computer. Spear-phishing webpages often resemble authentic pages on the victim's corporate intranet or externally hosted sites intended for legitimate activities (reviewing health insurance, employee benefits, etc.). In these cases, it can be difficult to distinguish between legitimate links and malicious copies.

With proper security awareness training, the recipient may find the email suspicious and forward the email to a security analyst to investigate. The analyst determines that this is a new phishing attempt and it would be useful for other organizations to be aware of this attack. What information should the analyst share and how?

There are several subtleties involved in this example. Simply sharing the threatening email is not enough. An analyst examining the mail will perform several actions, such as:

1. Look at the source IP address and determine the owner to see if the address is the origin of any previous attacks.
2. Check to see if there are any hyperlinks in the message text and, if there are, see if the target of the link is a known bad or suspicious domain.
3. Examine all attachments to see if these contain any hidden malware or executables.
4. Consider the text of the email to see what information is being requested, and whether the name of the sender matches the sender's legitimate email address.

It is likely that the majority of the information gleaned by the analyst could be extremely valuable to share with other organizations.

Turning Collaborative Intelligence to Action

Collaboration and aggregation of intelligence feeds is a necessary first step to combatting today's adversaries. However, without an analysis element, network defenders are still left sifting through largest volumes of indicators, potential threats, and other data that lacks sufficient context upon which to act. To provide this crucial layer of context, an analysis engine is required. The engine must deliver derived, relevant, and actionable threat intelligence based on observed incidents, case investigations, and social interaction in a collaborative community.

Once analyzed, the results can be disseminated to allow for action to be taken in an automated manner. Policies can be applied, along with consistency in what is shared, and confidentiality of the source. As the spear phishing campaign spreads to additional organizations, all could be better prepared by participating in an automated threat information exchange.

Conclusion

Threats are not just increasing in frequency, they are more targeted and customized than ever before. The collaboration of marketplace adversaries makes the problems much more complex and provides a constant stream of new attack vectors. In the face of these threats, most enterprises are, understandably, overwhelmed. The IT industry can no longer block against specific adversaries; it's now a marketplace of adversaries.

Sharing actionable threat intelligence in a secure, confidential, and timely manner is key. The benefits of such a security clearinghouse are widely recognized, yet no practical technology to support bi-directional automated security data sharing has proven to be successful.

Instead, organizations rely on emails, Web forums, and other advisories that must be manually processed and acted upon. Manually tracking and processing this data creates a significant challenge for IT security departments and doesn't scale well.

Sharing information about the source and environment of attacks allows us, as a community, to quickly isolate malicious or compromised hosts. In addition, information related to adversary attack patterns helps identify new tools and methods that can assist security research on new defense technology. Vendors like HP, with broad footprints in the security industry and leading security research, who can leverage vast amounts of real-time data, are best positioned to help solve this challenge.

A Security Research foundation

HP Security Research conducts innovative research and delivers intelligence to the full portfolio of HP Enterprise Security solutions, giving customers industry-leading protection against the latest threats. Security research publications and regular threat briefings complement the intelligence delivered through HP solutions and provide insight into the future of security and the most critical threats facing organizations today. Leading the company's security research agenda, HP Security Research leverages existing HP research groups, including HP DVLabs, HP Fortify Software Security Research, and manages the Zero Day Initiative (ZDI). Research areas of focus include vulnerability, malware, threat actor, and software security with a focus on the technologies, industries, and geographies most relevant today.

For the latest security research, visit hp.com/go/hpsr.

About Threat Central

Threat Central, developed with HP Labs, is a collaborative threat intelligence platform enriched by dynamic analysis and scoring.

Learn more at
hp.com/go/threatcentral

Sign up for updates
hp.com/go/getupdated

Share with colleagues

Rate this document

4AA4-6187ENW, June 2015, Rev. 1