Defining a new methodology for modeling and tracking compartmentalized threats

blog.talosintelligence.com/compartmentalized-threat-modeling/

May 13, 2025

By Edmund Brumaghin, Asheer Malhotra, Ashley Shen, Vitor Ventura

Tuesday, May 13, 2025 06:00

initial access broker

- In the evolving cyberthreat landscape, Cisco Talos is witnessing a significant shift towards compartmentalized attack kill chains, where distinct stages — such as initial compromise and subsequent exploitation — are executed by multiple threat actors. This trend complicates traditional threat modeling and actor profiling, as it requires understanding the intricate relationships and interactions between various groups, explained in the previous blog.

- The traditional Diamond Model of Intrusion Analysis' feature-centered approach (adversary, capability, infrastructure and victim) to pivoting can lead to inaccuracies when analyzing "compartmentalized" attack kill chains that involve multiple distinct threat actors. Without incorporating context of relationships, the model faces challenges in accurately profiling actors and constructing comprehensive threat models.
- We have identified several methods for analyzing compartmentalized attacks and propose an extended Diamond Model, which adds a "Relationship Layer" to enrich the context of the relationships between the four features.
- In a collaboration between Cisco Talos and The Vertex Project, a Synapse model update has just been published which introduces the *entity:relationship* providing modeling support to this methodology.
- We illustrate our investigative approach and application of the extended Diamond Model for effective pivoting by examining the ToyMaker campaign, where ToyMaker functioned as a financially-motivated initial access (FIA) group, handing over access to the Cactus ransomware group.

Impacts on defenders

The convergence of multiple threat actors operating within the same overall intrusion creates additional layers of obfuscation, making it difficult to differentiate the activities of one threat actor from another, or to identify when access has been handed off from one to the next. At each point where outsourcing occurs or access is handed off, the Diamond Model of the adversary changes. Likewise, the ability to leverage the output of kill chain analysis for the purpose of pivoting, clustering, and attribution becomes significantly more difficult as analysts may be forced to operate under the assumption that multiple actors are involved unless they can prove otherwise, where historically the opposite assumption was likely made.

Additionally, misattributing attacks due to tactics, techniques and procedures (TTPs) present in earlier stages of the intrusion may impact the way in which incident response or investigative activities are conducted post-compromise. They may also create uncertainty around the motivation(s) behind an attack or why an organization is being targeted in some cases.

Analysis processes and analytical models must be updated to reflect these new changes in the way that adversaries conduct intrusions, as existing methodologies often create more confusion than clarity.

Introduction to threat modeling

NIST SP 800-53 (Rev. 5) defines threat modeling as “a form of risk assessment that models aspects of the attack and defense sides of a logical entity, such as a piece of data, an application, a host, a system, or an environment.”

For many organizations, this involves evaluating their preventative, detective and corrective security controls from an adversarial perspective to identify deficiencies in their ability to prevent, detect or respond to threats based on specific tactics, techniques, and procedures (TTPs). For example, adversary emulation simulates an attack scenario and demonstrates how an organization could reasonably expect their security program to respond if a specific threat is encountered.

Intrusion analysis is the process of analyzing computer intrusion activity. This involves reconstructing intrusion attack timelines, analyzing forensic artifacts and identifying the scope and impact of activity. Intrusion analysis typically results in a better understanding of an attack or adversary, and may also result in the development of a model to reflect what is known about the threat. This model can then be used to support more effective detection content development and threat modeling activities in the future. The symbiotic relationship between intrusion analysis and threat modeling allows organizations to effectively incorporate new knowledge and information about threats and threat actors into their security programs to ensure continued effectiveness.

Over the past several years, different analytical models have been developed to assist with intrusion analysis and threat modeling that provide logical ways to organize contextual details about threats and threat actors so that they can be communicated and incorporated more effectively. Two of the most popular models are the Diamond Model and the Kill Chain Model.

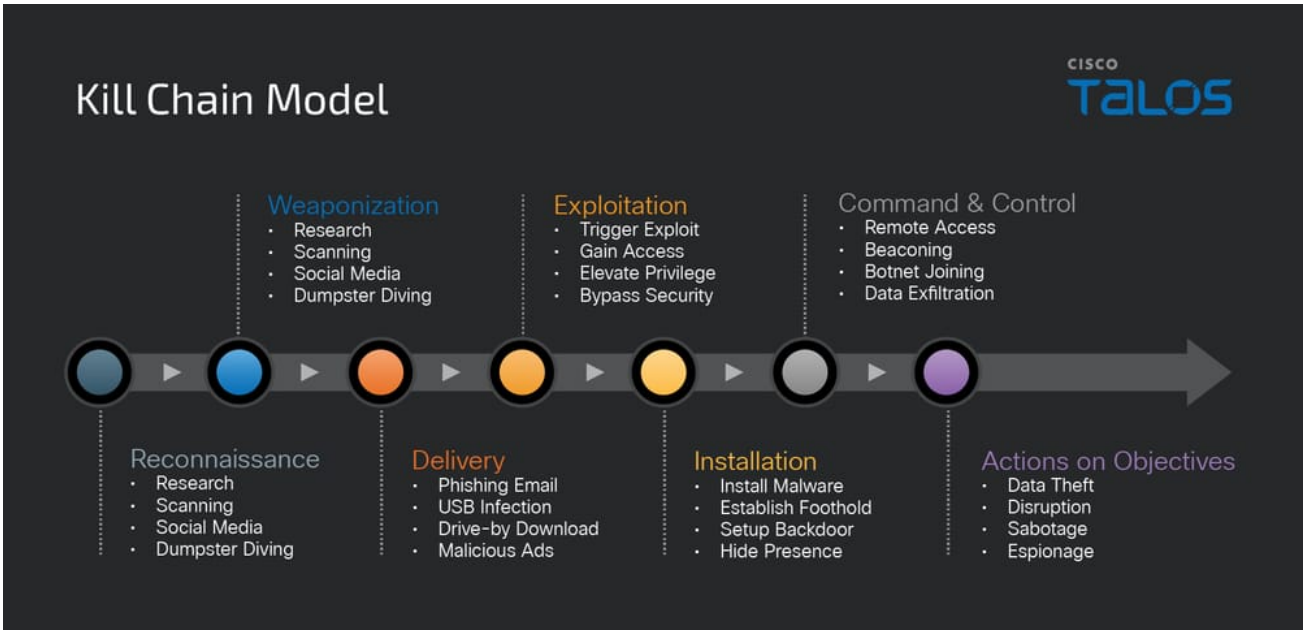

The Kill Chain Model shown above is typically used to break an intrusion down into distinct stages/phases so that the attack can be reconstructed and analyzed. This allows analysts to build a realistic model that reflects the TTPs and other characteristics present during the intrusion. This information can then be shared so that other organizations can determine whether their own security controls would be effective at combatting the same or similar intrusion(s) or whether they have encountered the same threat in the past.

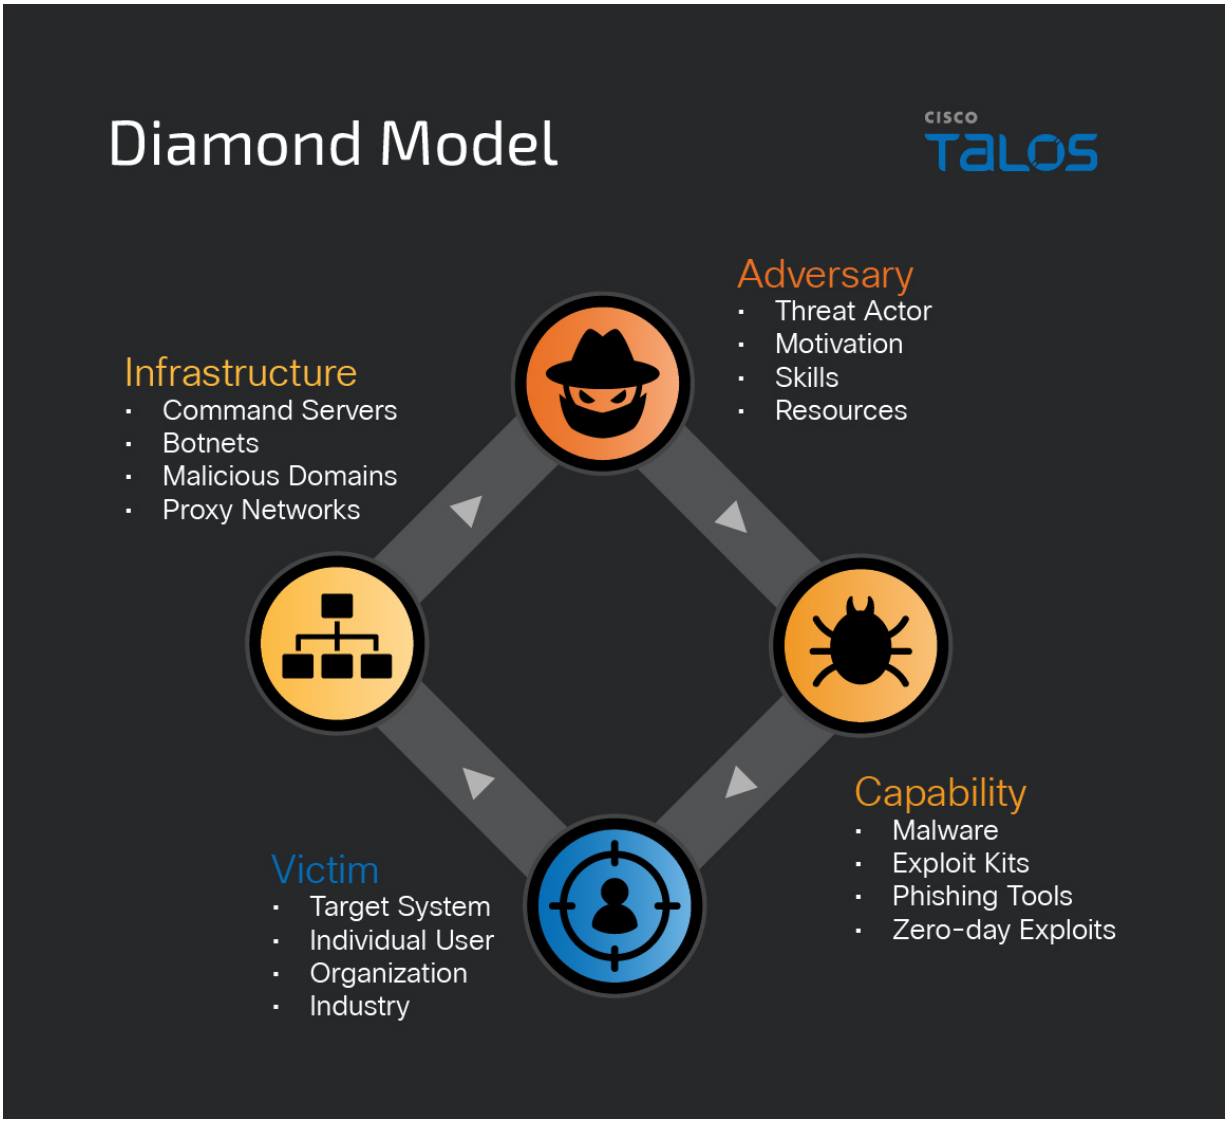

The Diamond Model, shown above, is commonly used across the industry for building a profile of a specific threat or threat actor. This model is developed by populating each quadrant based on information about an adversary's characteristics, capabilities, infrastructure tendencies and typical targeting/victimology. A fully populated diamond model creates an extensive profile of a given threat or threat actor.

It is important to note that an analysis may incorporate both (or other) models, and they are not mutually exclusive. There are also several other modeling frameworks that exist for similar purposes that are also often used in concert, such as the MITRE ATT&CK and D3FEND frameworks. For example, in some cases the information used to populate the

Diamond Model may be the result of kill chain analyses of multiple intrusions over time that are ultimately attributed to the same threat actor(s). By leveraging the output of multiple kill chain analyses, one can build a more comprehensive model that reflects changes to characteristics or TTPs associated with a threat actor being tracked over time as well as improve overall understanding of the nature of a given threat.

Challenges applying existing models to compartmentalized threats

One of the key strengths of the Diamond Model is its concept of “centered approaches” for analytic pivoting — including victim-, capability-, infrastructure- and adversary-centered methods of investigation. These approaches enable analysts to uncover new malicious activities and reveal how each facet of an intrusion across the Diamond’s four dimensions intersects with others. For instance, in the paper’s infrastructure-centered example, an analyst might begin with a single IP address seen during an intrusion, then pivot to the domain it resolves to, scrutinize WHOIS registration details, and discover additional domains or IPs registered by the same entity. Further examination may reveal malware connected to or distributed by those domains. In such scenarios, the Diamond Model’s systematic method of traversing from one node to another can rapidly expose an interconnected web of adversaries, capabilities, and victims.

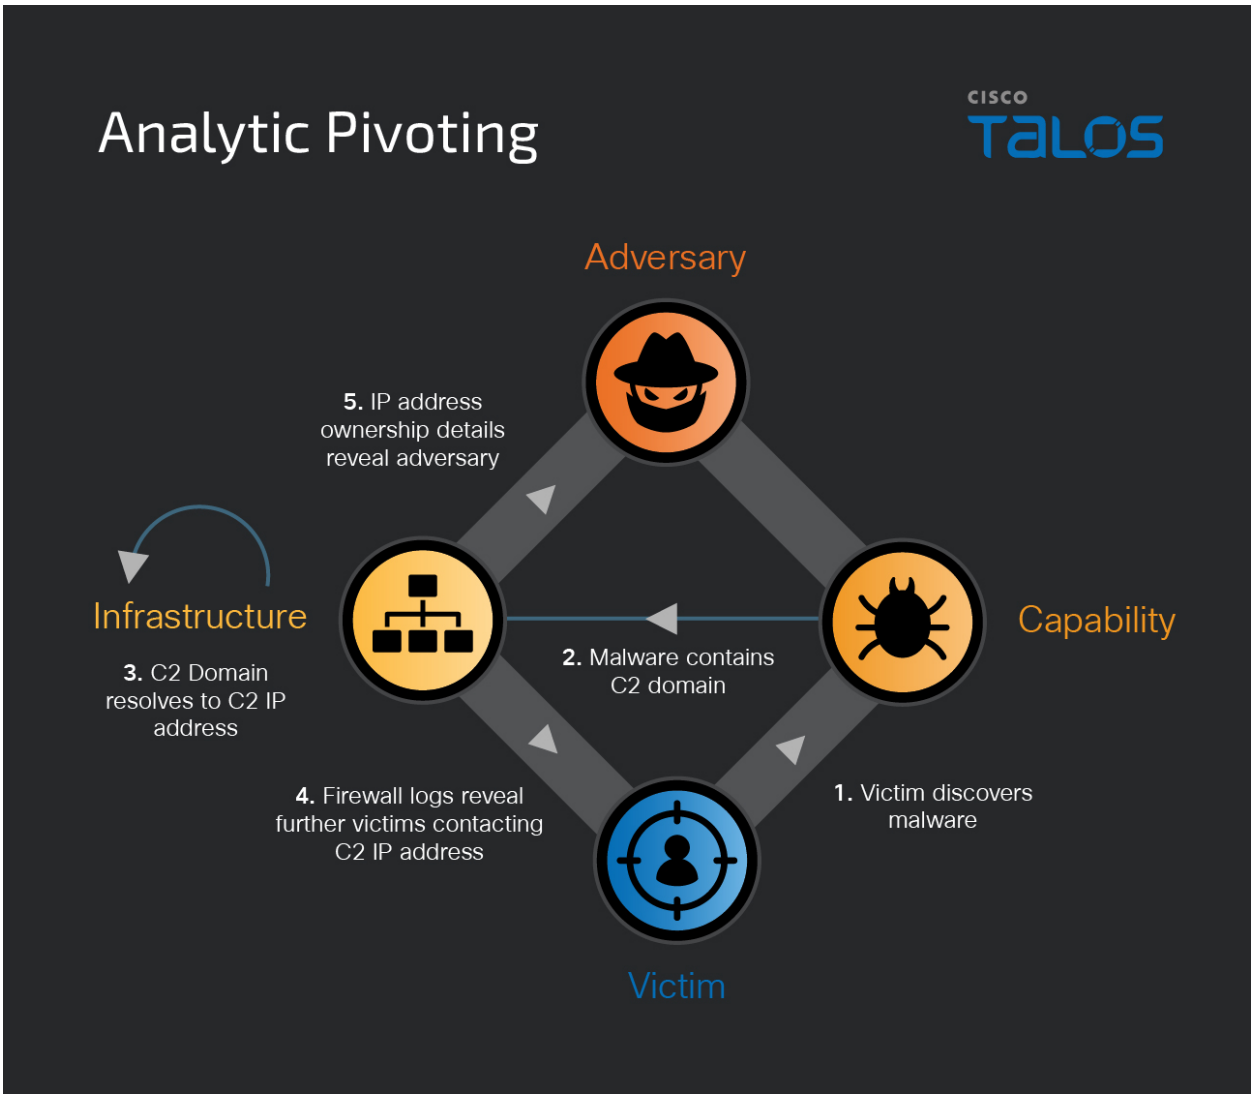

However, the original centered approach can introduce errors when dealing with a "compartmentalized" attack kill chain involving multiple distinct threat actors. In many cases, adversaries are now leveraging various relationships simultaneously while working towards their longer term mission objectives. This could include the outsourcing of tooling development, rental of infrastructure services for distribution or command and control (C2), or access-sharing agreements leveraged post-compromise to facilitate hand-off once initial access (IA), persistence or privilege escalation has been achieved. This compartmentalization has complicated many analytical activities including attribution, threat modeling, and intrusion analysis. Likewise, the modeling methodologies that were initially developed to combat intrusion operations in previous years no longer accurately reflect today's threat landscape.

To illustrate the complexity of compartmentalization, let's consider a hypothetical scenario that closely mirrors real-world events. In this scenario, four distinct threat actor groups are involved:

1. Actor A: A financially motivated threat actor aiming to profit by collecting logs from infostealer malware.
2. Actor B: A malware developer who creates and sells infostealer malware.
3. Actor C: A Traffic Distribution Service (TDS) provider.
4. Actor D: A ransomware group.

In this scenario, a financially-motivated threat actor (**Actor A**) who is seeking to infect victims with information-stealing malware to steal victims' sensitive information may outsource the development of their malware to **Actor B**. They may engage the developer directly or purchase it from a storefront. Likewise, the distribution of the malware itself is conducted by outsourcing it to **Actor C**, who operates a spam botnet or traffic distribution service (TDS) that is offered for rent for a usage-based fee. Once **Actor C** has successfully achieved code execution on a system, they may infect it with the malware they initially received from **Actor A**, who is charged "per-install."

Likewise, once **Actor A** has successfully performed enumeration of the environment, they identify that they were successful in gaining access to a high value target. Rather than simply focus on monetizing information-stealing malware logs, they choose to monetize their access to the exfiltrated data by selling it to **Actor D**, who then leverages that access to deploy ransomware and extort the victim.

In this hypothetical scenario, **Actor C**, who would be classified as a financially-motivated initial access (FIA) broker, may also be distributing multiple malware families at any given time and leverage traffic filtering to manage final payload delivery. They may even host these payloads on the same infrastructure. The nature of the business relationships described in this scenario are shown below.

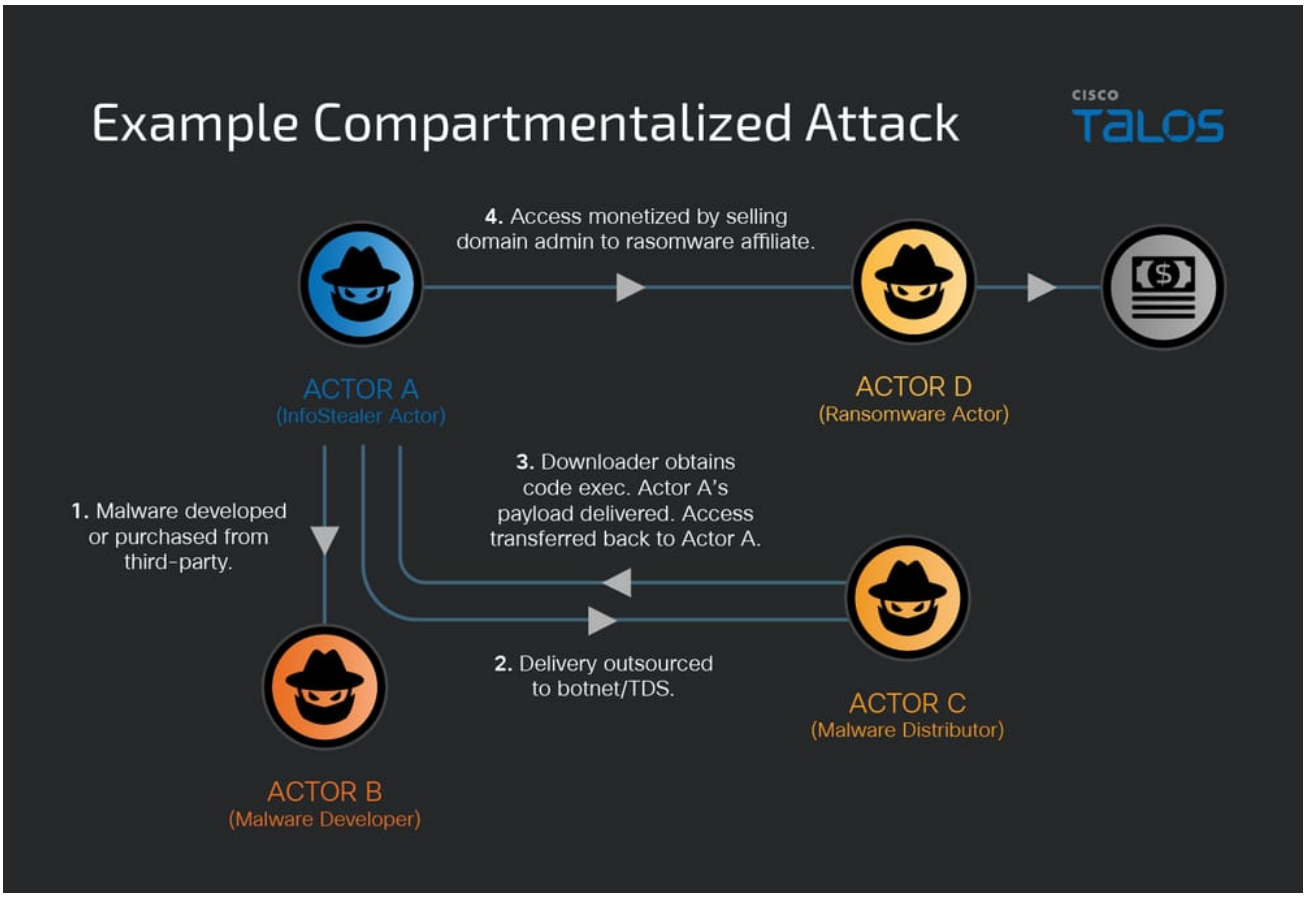

While this scenario covers a single attack, it highlights a situation where applying the traditional analytical models poses several challenges. For example, consider the infrastructure used by Actor C, the TDS provider. The infrastructure that facilitates malware distribution is not solely dedicated to Actor A's operations. This means that other malware found by pivoting the distribution infrastructure should not be considered as capabilities associated with Actor A. In addition, the malware's targets are highly associated with the Actor C's targeted network and should not be strongly considered as the motivation for the victimology of Actor A. In this compartmentalized scenario, the interconnected web of adversaries, capabilities and victims exposed by pivoting with the Diamond Model should not be associated with each other, as they originate from different threat actors and should not be modeled as part of a single threat actor profile.

In even more complex cases, a threat actor may choose to engage multiple distributors simultaneously or work with different distributors on a weekly basis depending on real-time pricing and service availability. A threat actor conducting ransomware operations may choose to procure access from several initial access brokers (IABs), each with their own characteristics, capabilities and motivations. Likewise, several otherwise unrelated threat actors operating in different capacities throughout the kill chain present complications when attempting to take the result of the analysis and incorporate it into existing attribution data or when attempting to identify overlaps with other clusters of malicious activity. Modeling the

IABs themselves also presents complications, as their characteristics and TTPs are often encountered in attacks where they may have only been operating within a subset of the overall phases of the intrusion.

State-sponsored or -aligned threat actors' campaigns have been documented using anonymization networks or residential proxies to hide their activities. This will create the same kind of activity overlap described by the usage of a TDS.

Extending the Diamond Model with the Relationship Layer

To extend the Diamond Model to include the complexities posed by compartmentalized attacks, we propose an extension to the original Diamond Model by integrating a "Relationship Layer." This additional layer is designed to contextualize the interactions between the four features (adversary, infrastructure, capability and victim) of individual diamonds representing distinct threat actors. By incorporating this layer, threat analysts can construct a nuanced understanding of compartmentalized contexts.

The Relationship Layer allows for the articulation of common relational dynamics such as "purchased from" to indicate a transactional association, "handover from" to reflect a transfer of operational control or resources, and "leaked from" to convey the use of leaked tools. Additionally, it describes the connections between adversarial groups, encompassing a variety of interactions such as "commercial relationship," "partnership agreements," "subcontracting arrangements," "shared operational goals," and more.

The integration of the Relationship Layer enables analysts to contextualize the interactions within the Diamond Model's four features, thereby enhancing their ability to perform logical pivoting and accurate attribution. This refinement offers a more sophisticated framework for analyzing modern, compartmentalized cyberthreats, providing a clearer representation of the complex web of relationships that characterize these operations.

Let's look at the scenario involving Actors A through D again. Figure 4 shows how we can use the extended Diamond Model to describe the relationships between entities involved in the intrusion activity:

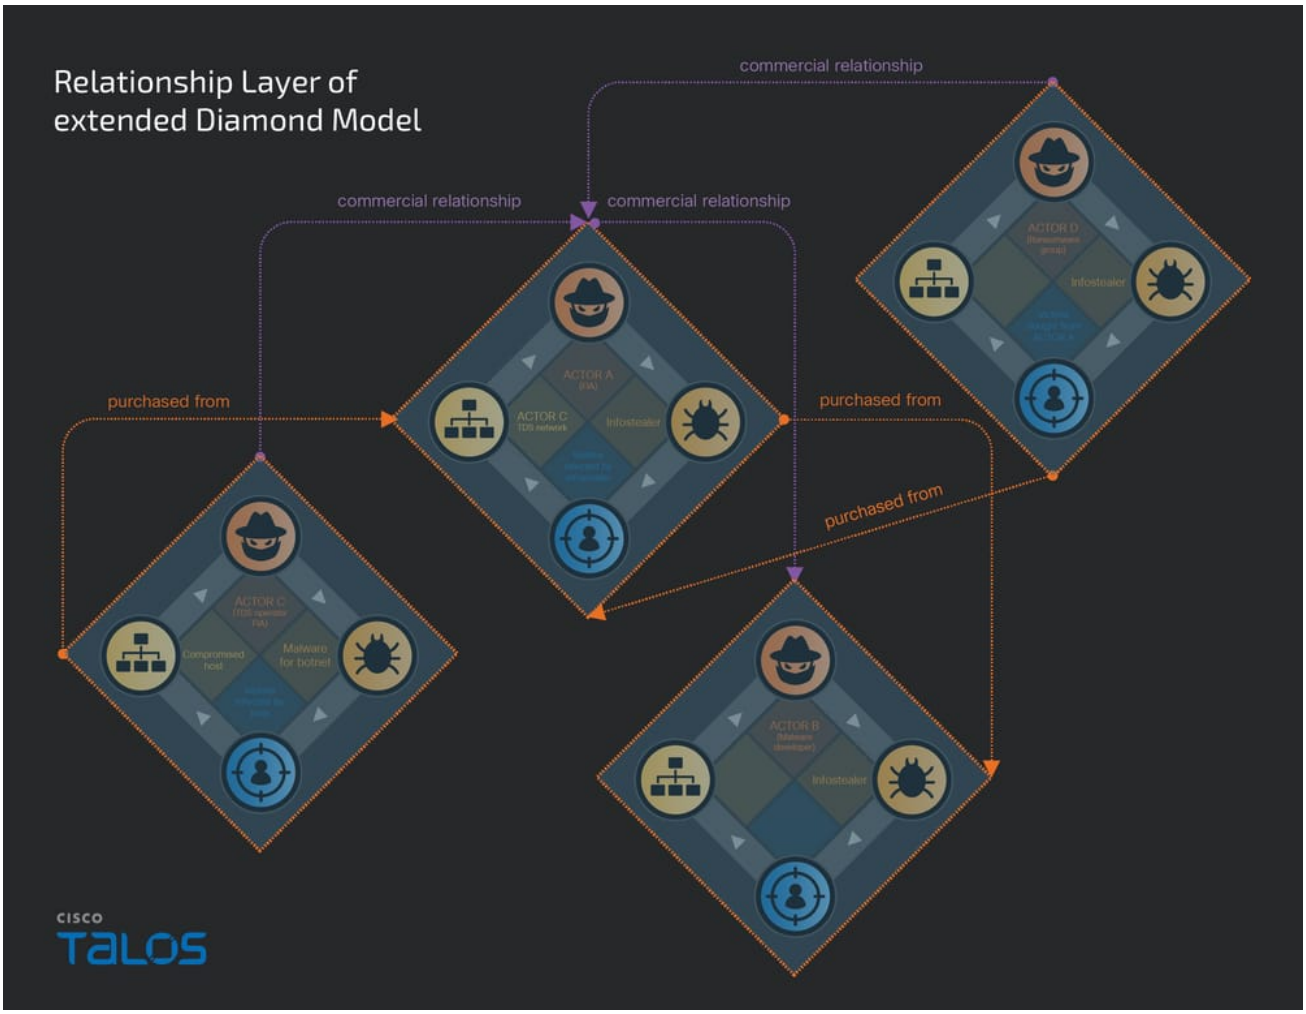

Each of the actors, A through D, possesses their own Diamond Model, reflecting their distinct roles as adversaries with unique capabilities, victims and infrastructures. We have extended each Diamond Model by integrating an additional Relationship Layer to illustrate the contextual relationships between these features. For instance, the infrastructure used by Actor A for Traffic Distribution Services (TDS) is linked to Actor C's infrastructure through a "purchased from" relationship. Consequently, when performing analytical pivoting, analysts should account for this relationship and not attribute all infostealers distributed via the TDS infrastructure solely to Actor A's capabilities. Similarly, the victims of those infostealers should not be automatically classified as Actor A's victims.

Another illustrative case involves the relationship between the victims of Actor A and Actor D. Actor D obtained initial access through a transaction with Actor A, denoted by the "purchased from" relationship within the Relationship Layer. This relationship offers analysts crucial context, allowing them to avoid attributing the tools used in the initial access phase to Actor D's capabilities.

The Relationship Layer also elucidates the connections between adversaries. On the graph, we denote these inter-adversary connections as "commercial relationships," providing additional context that aids in actor profiling. This extension understanding allows analysts

to discern the nature of interactions between threat actors, facilitating more accurate and insightful profiling efforts.

Integrating the Relationship Layer with the Cyber Kill Chain

The Cyber Kill Chain framework serves as a structured approach to analyzing cyberattacks, enabling security professionals to break down intrusions into discrete, sequential stages — from initial reconnaissance to actions on objectives. By organizing attacks in this manner, analysts can pinpoint attacker behaviors, anticipate adversary actions and develop targeted mitigation strategies, significantly enhancing overall threat intelligence.

Integrating the extended Diamond Model into the Cyber Kill Chain framework offers a more comprehensive view of compartmentalized campaigns by illustrating how each adversary contributes to different stages of an attack. This combined perspective enhances understanding by mapping out the intricate web of relationships among multiple threat actors, thereby providing a clearer picture of how resources, capabilities and infrastructure are shared or transferred throughout an attack's lifecycle. Figure 5 illustrates the integration of the extended Diamond Model with the Cyber Kill Chain using the Actor A–D example.

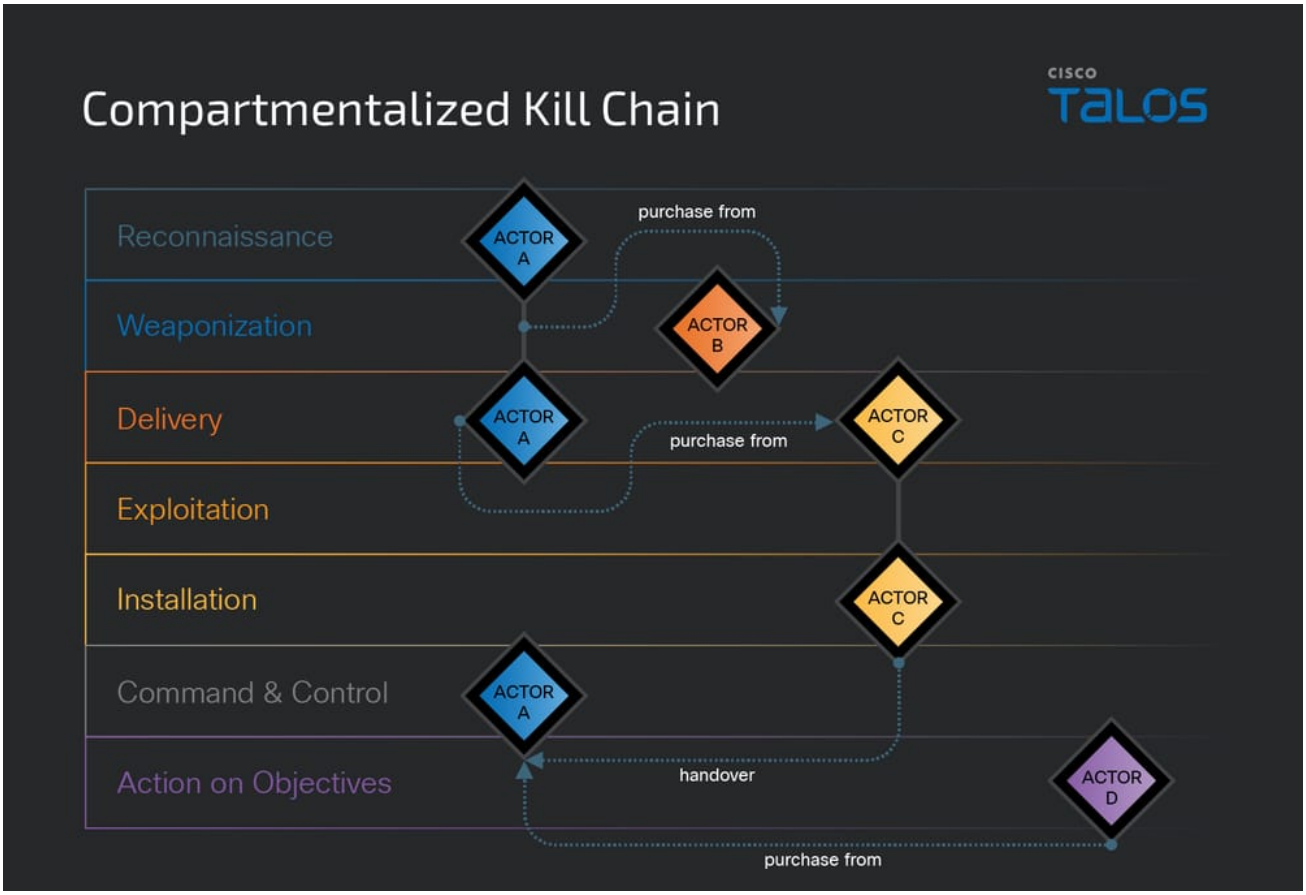

The example above demonstrates the distinct roles that each adversary assumes at various stages of the kill chain in a hypothetical campaign. In this scenario, the victim is initially compromised by an infostealer, which Actor A acquired from Actor B, and subsequently

faces a ransomware attack orchestrated by Actor D. To further enrich the analysis, we highlight the "handover" relationship between Actor C and Actor A, emphasizing its significance as both actors' activities manifest within the targeted environment. This approach provides a more comprehensive view of the attack flow, allowing for a deeper understanding of how adversarial interactions and transitions unfold throughout the campaign.

This enriched view not only clarifies attacker tradecraft but also bolsters actor profiling and attribution efforts. By aligning specific tactics and resources with the threat groups deploying them, analysts can more accurately trace operations back to their origins. This approach also provides insights into adversary motivations, allowing defenders to tailor their response strategies effectively. For instance, understanding that an IAB is financially motivated might suggest a lower immediate threat to certain targets, while recognizing that access has been sold to a state-sponsored actor would escalate the priority of the threat response.

Identifying compartmentalized attacks

Identifying compartmentalization within the scope of an intrusion typically involves trying to determine where positive control is transferred between adversaries either pre- or post-compromise. It is essential to identify compartmentalization as this will significantly impact the overall understanding of the adversar(ies) and the capabilities available to them. Indicators of collaboration among distinct threat actors can vary significantly depending on the context and the phase of activity, and these can be categorized based on whether the actions occur before or after the compromise of a system or environment. It is important to note that while there are several examples listed in the following sections, compartmentalization can and does look different across intrusions and these are by no means comprehensive. Likewise, while the below elements are useful indicators that an analyst should investigate possible transfer of access, they are not necessarily indicative that a handoff has occurred. As more of these elements are encountered and evidence collected, an analyst may be able to strengthen their assessment that compartmentalization has occurred.

Pre-compromise

In the early stages of an intrusion, compartmentalization can often be identified by observing how tooling has been sourced, how malicious content is being delivered to potential victims and the initial/early execution flow of malicious components in the case that code execution has been achieved.

This stage may also be completely independent. In situations where a state-sponsored group is tasked with espionage operation, it may pass on the access to a ransomware group, making the state-sponsored group an IAG. It is not guaranteed that the ransomware

group is aware of the nature of its IAG, but just by doing its activity it will fulfill the state-sponsored group objective of making incident analysis and attribution complex.

Shared tooling

While many of the indicators associated with the use of tooling are often identified in later stages of an intrusion, we characterize this compartmentalization as occurring pre-compromise as development and procurement activities must generally occur before the campaign is launched. It is often useful to identify if the threat actor procured tooling from third parties. This may involve identifying key characteristics of the malicious components being analyzed and searching/monitoring hacking forums and darknet marketplaces (DNMs) to identify whether a seller is advertising a capability matching the one used in the intrusion. Likewise, malware that has historically been used by one threat actor may be transferred to another threat actor, either on purpose or inadvertently in the case of source code leaks. In either case, analysis of contextual information surrounding the use of the tooling can help analysts identify when the tooling doesn't match the threat actors' known TTPs.

Shared delivery infrastructure

In the case of email-based delivery, analysis of the infrastructure used to send malicious emails, the content of the message, and the infrastructure used for hosting and delivering payloads may indicate that delivery has been outsourced in some capacity. Likewise, in the case of malvertising campaigns, analysis of the ad campaigns, traffic distribution infrastructure and gating methodologies may suggest the same. In many cases the infrastructure used is often observed distributing multiple distinct, otherwise unrelated malware families over a short period of time as the threat actor operating the delivery infrastructure may conduct business with multiple entities at any point in time. Analyzing activity associated with this infrastructure before, during, and after the intrusion may inform the analysis of whether compartmentalization has occurred.

Shared droppers/downloaders

When analyzing an intrusion, there is often a point at which code execution is achieved. This may be the point in which a malicious script-based component is delivered and executed by a victim. In many cases, these function as downloaders and are solely responsible for retrieving or extracting and executing follow-on payloads that allow an adversary to expand their ability to operate in an environment. Analysis of the dropper/downloader mechanisms used may identify cases where the same mechanism is used to deliver unrelated threats over time, indicating that delivery may have been outsourced. We have categorized this activity as "pre-compromise" to further differentiate it from handoffs that may occur later in the intrusion, once persistence has been achieved, etc.

Post-compromise

In addition to the aforementioned types of compartmentalization that often occur early in an intrusion, there is another set of handoffs that may occur once an adversary has achieved compromise. These are typically used to transfer control of access from one party to another and may be performed for a variety of purposes, as described in our previous blog. This activity can often be identified by analyzing handoff behaviors, the motivation of the threat actors involved, and monitoring for typical indicators that an IAB is involved.

Handoff behaviors

In some cases, information can be collected related to the amount of time that has occurred between an IAB obtaining access to the environment, and the beginning of follow-on activity. This may include an IAB gaining access, establishing persistence, collecting information from the environment and exfiltrating that to adversary C2. Following this initial activity, the infection may conduct very little malicious activity aside from periodic C2 polling occurring on the system for an extended period of time. After an extended period, additional malicious components may be delivered that establish new C2 connections and new activity may be observed. This type of pattern is indicative that a handoff of access may have occurred and should be investigated further. Similarly, analysis of the behaviors of the threat actor before and after this handoff may strengthen or weaken an assessment as completely different TTPs may be observed between the threat actors involved.

The race to domain admin

Another set of characteristics that may strengthen an assessment that handoff has occurred is by analyzing the series of actions taken once access has been gained. In the case of FIA, for instance, we often observe repeatable processes for attempting to gain domain administrator access as quickly as possible. This makes the access more lucrative for the IAG and more seamlessly enables the deployment of additional malware components, such as ransomware. An FIA group may quickly progress from initial access to domain administrator access in a short period of time with little to no effort spent on identifying high-value targets in the environment. Once domain administrator access has been gained the intrusion activity may stop while the threat actor attempts to monetize that access and facilitate handoff to the threat actor who ultimately purchases it. SIA groups on the other hand, may take a more steady and stealth oriented approach, to conduct reconnaissance and proliferate throughout the victim enterprise without being detected. In many instances an SIA group might conduct initial exfiltration of restricted data, before handing access off to the secondary threat actor.

Dark web tracking

Monitoring hacking forums and darknet marketplaces can be extremely valuable for identifying when an IAB is involved in an intrusion. Since FIA brokers are primarily focused on achieving the maximum profit as quickly as possible, they will often post advertisements for access to environments that they have achieved. In many cases these advertisements include generic information about the company/organization involved such as size (number of employees), rounded financial information based on publicly available sources such as quarterly filings, industry, etc. Locating advertisements that match the profile of the victim of an intrusion can strengthen an assessment that an IAB is involved and provide additional intelligence collection avenues that may be pursued further to collect additional information about the IAB involved, who they typically work with, and more.

C2 analysis

Analysis of C2 infrastructure involved throughout the intrusion presents another opportunity for identifying any handoffs that have occurred. As previously mentioned, in some cases the handoff is performed by delivering a new payload and establishing a new C2 connection with another threat actor's infrastructure. In the case of frameworks, analysis of the server logs can provide additional information where the same server has been used to administer multiple victims. Administrative panels used to manage malware infections are often useful for informing analysis related to the nature of threat actors involved and the business models they are working within. Some admin panels may be explicitly built for the purpose of facilitating handoffs, RaaS and C2aaS platforms being examples of this.

Case Study: ToyMaker

During the course of performing threat hunting and incident response, Cisco Talos sometimes encounters scenarios where compartmentalized operations involve multiple attackers participating in the same attack kill chain. Using the ToyMaker campaign as an example, we demonstrate how we identified the participation of various attackers during our investigation and utilized the extended Diamond Model to clarify the distinct activities and roles of these attackers across different stages of the attack kill chain.

APT, Cactus or FIA?

Talos investigated the ToyMaker campaign in 2023. The attackers conducted operations for six consecutive days, during which they compromised a server of the victim organization, exfiltrated credentials and deployed the proprietary LAGTOY backdoor. We consider this "first wave" post-compromise activity. Since we did not find any common financial crime malware in this attack, and the attackers used their proprietary tools and C2 infrastructures, we considered the possibility that it might be the activity of an APT group. However, the TTPs and indicators of compromise (IOCs) did not overlap with previously observed campaigns, so we did not attribute the campaign early in the investigation.

However, during the investigation, Talos identified TTPs and hands-on-keyboard activity consistent with Cactus ransomware activity appearing in the victim's network almost 3 weeks after the initial compromise. We consider this the “second wave” of malicious activity. After using various tools for lateral movement within the network, the attackers launched a ransomware attack within a matter of days. At this point, Talos started a more in-depth investigation, including exploring the connections and disparities between the ransomware attack and the initial access. We formulated several hypotheses at this point:

- **Hypothesis A**: Both the initial compromise and subsequent activities were conducted by Cactus ransomware, and therefore LAGTOY might be a tool exclusively used by Cactus.
- **Hypothesis B**: The initial access might have been carried out by a different attack group and have no relation to Cactus's activities.
- **Hypothesis C**: The initial access might have been carried out by a different attack group, but there is some connection to Cactus.

Hypothesis A was the most intuitive assumption at the beginning of the investigation. However, as the investigation progressed, Talos made the following observations:

- **Initial access activity removed the created user account before the end of activity:** Before the actions following the initial access activity ceased, the attackers deleted the user account they had created.
- **Differences in TTPs:** Variations in TTPs were observed between the two attack traces, either through differing approaches to similar TTPs or entirely distinct TTPs. For instance, the operators conducting initial access relied on PuTTY for credential exfiltration, while the secondary activity employed Secure Shell (SSH) alongside other tools. In terms of file packaging, the second wave utilized parameters that preserved file paths (-spf), a method not seen in the first set of actions. Furthermore, the second wave predominantly involved off-the-shelf tools, whereas the first wave featured bespoke tools unique to the attackers.
- **No tools and IoC overlapping:** We found no common tools and shared infrastructure between the two waves of malicious activity.
- **No use of LAGTOY:** We observed that although the first wave deployed LAGTOY, it was never used throughout the course of the intrusion. Why would a threat actor deploy a custom-made malware immediately after initial compromise but never use it? It is possible that LAGTOY might have been designated as a last resort access channel, if the attackers’ access through compromised credentials was blocked. It is also likely that LAGTOY wasn’t used because it was never meant to be used in the intrusion going forward, i.e. LAGTOY was deployed by a distinct Initial Access Threat Actor, different from Cactus. Furthermore, we had no evidence of Cactus developing and using LAGOTY in their operations. Our assessment was now leaning towards **Hypothesis B:** The initial access might have been carried out by a different attack group and have no relation to Cactus's activities.

- **Time gap between the first and second waves:** There was approximately a gap of 3 weeks with no observed attack activity before the second wave of attacks began. For big-game double extortion threat actors, speed is paramount. A successful initial compromise must be capitalized by performing rapid recon, endpoint and file enumeration, data exfiltration and ransomware deployment. For such operations that tend to focus on a blitz, it is abnormal to see a gap of weeks with lulls in activity. Therefore, we must consider the possibility that there may have been a handoff of access between two distinct threat actors conducting the first and second wave of attacks. Furthermore, a gap of 3 weeks suggests that the first threat actor did not have a secondary actor already aligned/available for immediate access; they had to find Cactus. Talos' assessment was now leaning towards **Hypothesis C**: The initial access might have been carried out by a different attack group, but there is some connection to Cactus.

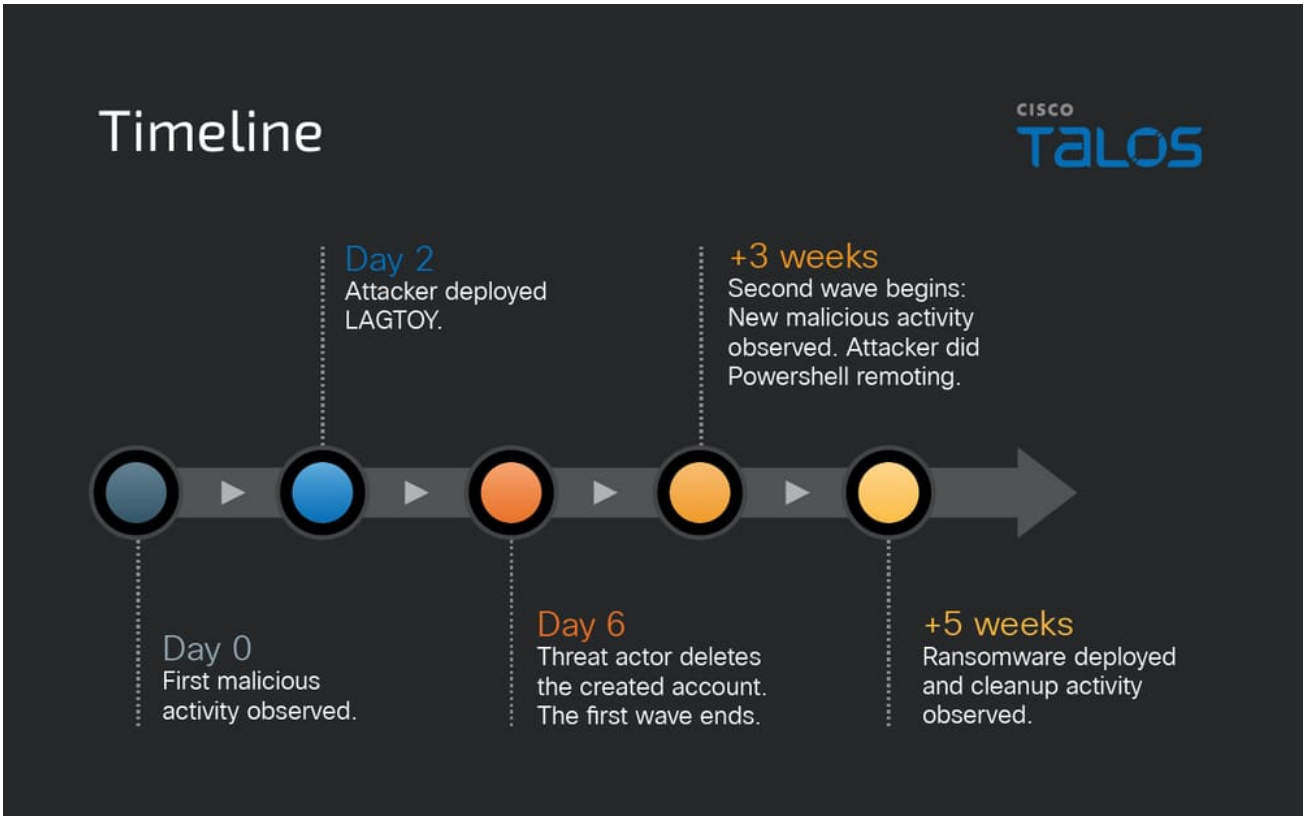

Shared credentials: Within the first six days of activity, we observed credential harvesting and exfiltration. Three weeks later, the second wave began which we attributed to Cactus. This second wave was kickstarted using the same credentials stolen in the first wave. Therefore, there was indeed a connection between the two waves of activity: the shared stolen credentials.

The totality of patterns and abnormalities collected during our research shifted our assessments toward the hypothesis involving an initial access group, leading us to reanalyze the LAGTOY tool used in the first wave of activities conducted post compromise.

We discovered that this backdoor is the same as HOLERUN, which Mandiant reported as being used by UNC961. This finding, combined with the previous public reporting and observations, allows us to confirm that the attack involved two distinct attacker groups (ToyMaker aka UNC961, and Cactus).

Mandiant's public reporting noted that UNC961's intrusion activities often preceded the deployment of Maze and Egregor ransomware by distinct follow-on actors. While Egregor is considered a direct successor to Maze, there is no evidence indicating any connection to Cactus. In the campaign we investigated, Cactus used compromised credentials from the first wave of attacks on the victim's machine. Based on these findings, Talos assesses with high confidence that ToyMaker provided initial access for the Cactus group. Given ToyMaker's focus on financial gain and their history of selling initial access to ransomware groups, we classify them as an FIA group.

Leveraging the extended Diamond Model for further analysis and defensive strategy

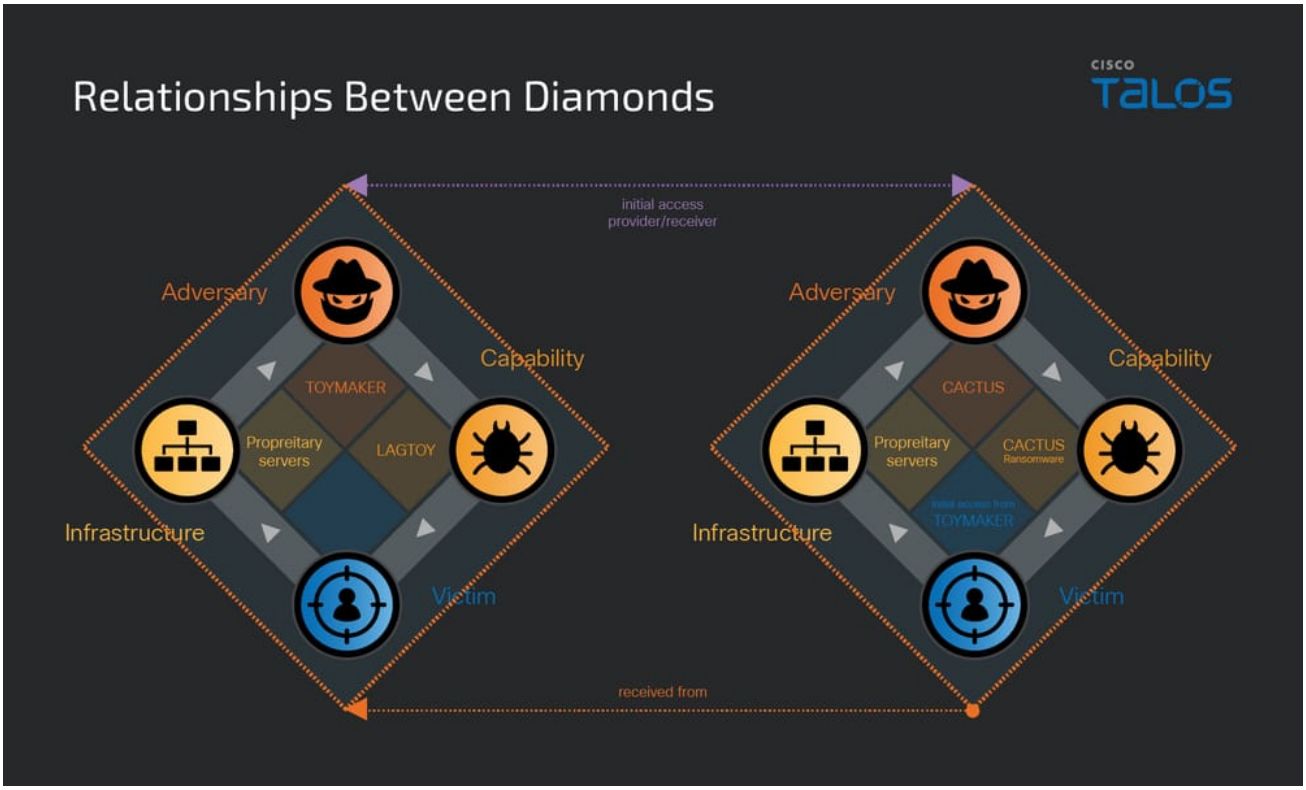

Building on the analysis and context provided, the extended Diamond Model allows Talos to effectively represent the threat actors involved in this campaign, highlighting the intricacies of their collaborative relationships. In Figure 6, we utilize two distinct diamonds to symbolize the ToyMaker group and the Cactus ransomware group. The Relationship Layer plays a crucial role in delineating the connections between ToyMaker's victims and Cactus' victims, as well as illustrating the initial access provider-receiver dynamics.

These relationships underscore the importance of carefully reviewing and investigating any capabilities and infrastructure indicators identified on the victim's machine associated with either threat actor. For example, the hosts infected by LAGTOY are potentially at risk of ransomware attacks, or tools discovered on Cactus' victims might be from LAGTOY or potentially other initial access groups.

We can also leverage the relationship information provided by the extended Diamond Model to identify additional potential victims of Cactus ransomware by hunting for hosts infected with the LAGTOY backdoor. Similarly, examining victims associated with ToyMaker can lead to discovering other ransomware attack victims. For defenders, this relationship data is crucial for prioritizing detection efforts and ensuring that the activities of ToyMaker and other initial access groups are not overlooked, as they can serve as precursors to further attacks. By maintaining vigilance and focusing on these initial access indicators, security teams can proactively identify and mitigate threats before they escalate into full-blown ransomware incidents.

© Cisco Systems, Inc. and/or its affiliates. All rights reserved. View our Privacy Policy.