

The Diamond Model of Intrusion Analysis

A Summary

By
Sergio Caltagirone

This document is not a reference guide to the Diamond Model. See technical report for official reference and complete details.

Why the Diamond Model Matters

The Diamond Model, for the first time, accurately details the fundamental aspects of all malicious activity as well as the core analytic concepts used to discover, develop, track, group, and ultimately counter both the activity and the adversary. The model emerged in 2006 by senior analysts asking the simple question, “How do we do our work?”

Unfortunately, it required seven years of thought, implementation, and refinement to complete the model. This delay is primarily because the intrusion analysis discipline has long been regarded as an art – to be learned and practiced, rather than a science – to be studied and refined. It is a discipline that prizes and studies analytic outcomes far more than understanding the processes and principles used to those achieve those outcomes. This approach has held analysis back from identifying first principles and foundational concepts. It frustrated the development of new tradecraft and a more complete understanding of malicious activity. This restriction had further implications slowing the evolution of threat mitigation which relies on efficient, effective, and accurate analysis.

The Diamond Model begins to address these challenges by applying scientific rigor to the discipline. With the Diamond, new and more effective mitigation strategies can be developed that increase the cost on the adversary while reducing the cost to the defender. It integrates traditional information assurance strategies and cyber threat intelligence seamlessly. It increases analytic efficiency and effectiveness by highlighting analytic opportunities and intelligence gaps. It achieves measurability, testability, and repeatability in the analytic process increasing analytic accuracy. It establishes the foundational concepts for cyber ontologies, taxonomies, and cyber threat intelligence sharing protocols. It adds context and relationships to previously flat and isolated indicators.

The Diamond is unique. It is at the same time formal and informal, simple and complex. It is a simple and informal enough to be used by analysts daily in pursuit of the adversary. Yet, it is complex and formal enough to establish a mathematical framework upon which to apply advanced concepts such as clustering, classification, game theory, graph analysis, and others.

Most importantly, the Diamond Model is just the beginning of the story as it can, and should, be critiqued and refined to improve the discipline while discovering previously unknown benefits and applications.

The Diamond Model of Intrusion Analysis

The Diamond Model's atomic element is the **event**. The event describes the four core features present in every malicious event: that for every intrusion event there exists an *adversary* taking a step towards an intended goal by using a *capability* over *infrastructure* against a *victim* producing a result.

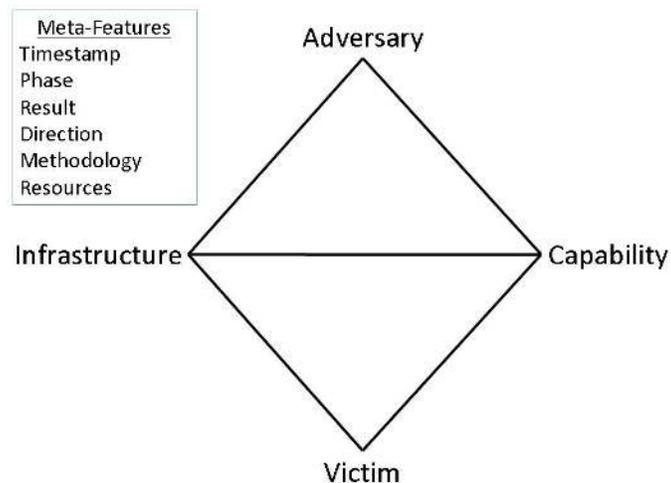


Figure 1. The Diamond Event

These core features are connected by edges which define a unique relationship between each of the features. This graph is now organized to resemble a diamond, hence the name. See Figure 1. The Diamond Event illustrating the Diamond features and their relationships.

The relationship/vertices between features are based on analytic pivoting and how from any point on the Diamond, an analyst can possibly reach the other connected points (given data/visibility/etc.). For example, from the victim an analyst can 'see' the capabilities being used against the victim over which infrastructure. From the infrastructure an analyst can 'see' the capabilities being used over the infrastructure and to which victims; additionally from the infrastructure, the analyst could potentially 'see' the adversary controlling the infrastructure. Moving from one feature of the Diamond to the others is called *analytic pivoting* and is a core analytic use of the model allowing analysts to maximize opportunities and clarify intelligence gaps.

The Diamond also defines additional meta-features of an event such as: *timestamp*, *phase* (e.g., reconnaissance, weaponization, exploitation), *result* (e.g., success, failure, confidentiality compromised, integrity compromised), *direction* (e.g., from or to the victim, bidirectional), *methodology* (e.g., spear-phishing, denial of service attack), and the external *resources* necessary to successfully complete the event activity (e.g., the target's email address or IP address, the vulnerability to exploit). These are included in the model as generally useful meta-features but are not core features and can be removed and augmented as necessary.

The Diamond is purposefully generic to enable each implementing organization the freedom to define, tailor, and augment meta-features as necessary. Furthermore, each core and meta-feature can be expanded and composed of many sub-features, which in themselves may be composed of sub-features allowing the model to become as complex and rich as necessary to describe the activity. For example: the victim may be defined by

the victim organization, the IP address of the victim asset being targeted, the hostname of the victim asset, etc. This malleable aspect makes the model useful in describing malicious intrusion activity in any organization. Yet, it is not so generic that it loses its key properties.

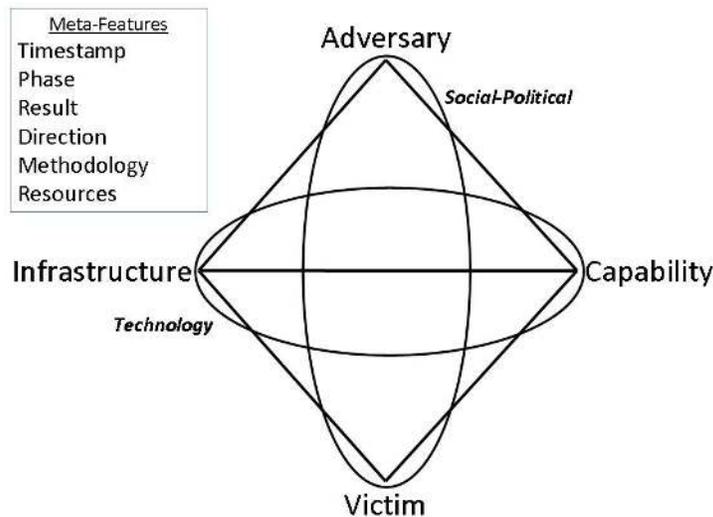


Figure 2. Extended Diamond Model

Domain Name System (DNS) to determine its command-and-control point, then DNS is part of the technology meta-feature.

The social-political meta-feature describes the always existing, and sometimes enduring, relationship between adversary and victim. It describes that there are underlying needs, aspirations, and intent behind every malicious activity – and that the victim plays a unique role in that relationship. Analytically, the Diamond allows concepts from criminology and victimology to be applied to intrusion analysis allowing one to understand the reason a victim was chosen, the value the victim brings to the adversary, and ultimately how that relationship can be influenced and manipulated to enhance mitigation. Particularly, this highlights the existence of a *shared threat space* where two or more victims satisfy the needs of one or more of the same adversaries. The existence of this space illustrates that the sharing of threat intelligence is more lucrative with those most likely to be impacted by a similar adversary as well as enables the members of the space to potentially forecast and predict future malicious activity.

As mentioned previously, the Diamond inherently supports analytic pivoting. This concept is further extended into the ‘centered’ approaches. These six approaches, from the four core features of the Diamond plus the technological and social-political meta-features, enumerate all of the potential methods to discover cyber threats. For example, one could use the infrastructure-centered approach to exploit the features of malicious infrastructure to discover new infrastructure. Or, one could use the victim-centered approach to carefully watch a victim of high interest to a likely adversary to discovery previously unknown aspects of their activity.

The Diamond is further expanded by two additional meta-features which define the technology and social-political meta-features. See Figure 2. Extended Diamond Model illustrating these additional meta-features.

The technology meta-feature connects the infrastructure and capability and describes the technology enabling the infrastructure and capabilities

to interact effectively. For example, if malware uses the

Per the first principles identified by the Diamond Model, adversaries operate in multiple phases. It requires a minimum of two or more events to cause a malicious outcome (i.e., at a minimum an adversary must conduct target selection followed by a malicious action which are two steps). Once activity is discovered and events have been

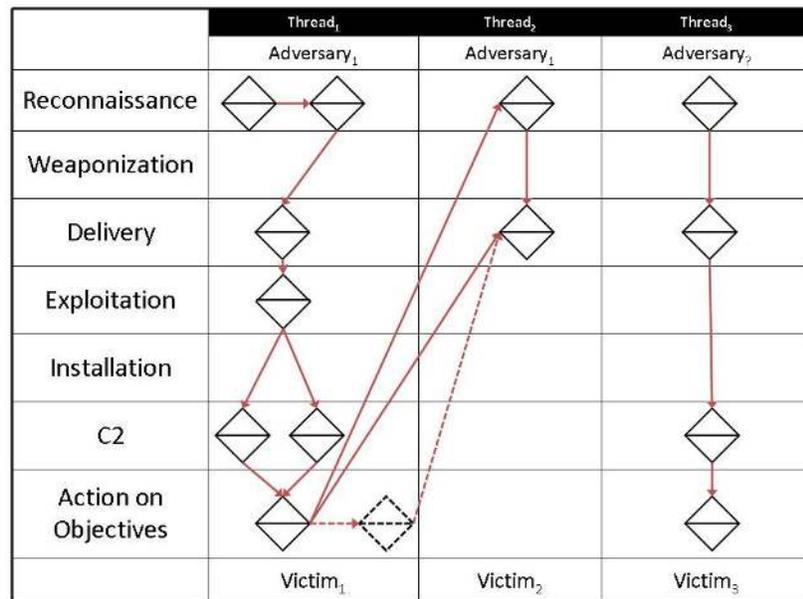


Figure 3. Activity Thread Example

characterized and analyzed, they are

ordered by the phases of malicious activity and linked by their causal relationship into threads. These are called **activity threads**. Threads not only span vertically along a single adversary-victim pair, but horizontally as adversaries take advantage of knowledge and access gained in one operation to enable other operations.

Figure 3. Activity Thread Example illustrates an adversary's operations against two victims as well as another unknown adversary's operations against a third unrelated victim. Furthermore, the dashed elements illustrate the ability for analysts to integrate hypotheses that can then be further tested or supported with additional evidence gathering. This organization of knowledge is useful in many ways, including: the identification of knowledge gaps and adversary campaigns, as well as hypothesis generation and documentation. Sub-graphs of these threads are called **adversary processes** which can be useful later to group and classify activity based on process rather than single indicators.

While activity threads organize the intelligence of adversary operations, attack graphs (i.e., attack trees) have historically been used by information assurance and information security groups to postulate all exploit paths to an asset. This methodology has been used to then plan defense decisions based on the cost of a defensive action, the number of exploit paths it covers, and the value of the asset (amongst other variables). However, this method has generally not survived either adversary contact or red-teaming because it requires omnipotence of the threat and vulnerability space, or a degree of exhaustion over that space to give oneself a reasonable level of assurance – albeit likely misplaced at times.

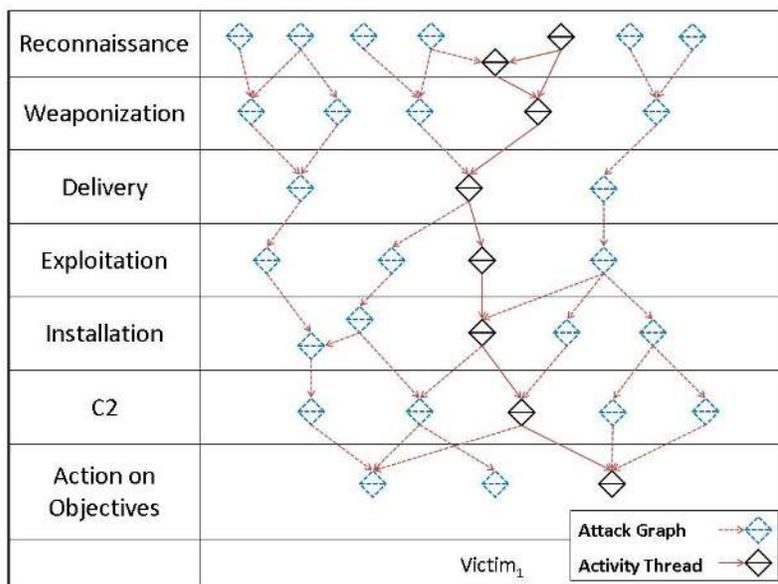


Figure 4. Activity-Attack Graph Example

The Diamond Model addresses this and integrates activity threads and attack graphs into a new structure called the **activity-attack graph**. This new graph allows defenders to now take intelligence of adversary operations into account during planning not only to sever existing operations paths, but also to potentially predict future paths based on adversary preference.

For example, given Figure 4. Activity-Attack Graph, if a defender severed the known activity thread delivery mechanism, and the adversary wanted to remain persistent on the victim, then the attack graph illustrates that there are two alternative delivery routes. Therefore, it may be prudent to not only take action to block the adversary’s currently known delivery mechanism, but also the other potential delivery mechanism thereby providing proactive mitigation and possibly pre-empting the adversary’s next maneuver.

While activity threads and activity-attack graphs give analysts and defenders a view of the adversary’s operations and processes, the next question is to identify adversary campaigns and group like activity to answer high-order analytic problems (e.g., such as identifying common capability developers, establishing pattern-of-life timelines, conducting longitudinal victimology to deduce intent and attribution).

Groups of common/similar malicious events, adversary processes, and threads are called **activity groups**. Analysts traditionally form activity groups to identify a common adversary using similarities in infrastructure and capabilities; but, the concept is inherently flexible and extends to include any grouping based on similarities. Furthermore, activity groups themselves can be hierarchical and organized into **activity group families** more accurately modeling mature organizations (e.g., organized crime) behind some cyber threat activity.

There are six steps to the process of activity groups:

1. **Analytic Problem Definition:** The analytic problem to solve by grouping is defined.

2. **Feature Selection:** The set of features on which to measure similarity between events and threads is defined.
3. **Creation:** The activity group is created by clustering similar events, threads, and processes using the features in the feature vector to solve the analytic problem.
4. **Growth:** The activity group is grown by classifying new events, threads, and processes into the existing groups as they are discovered.
5. **Analysis:** The activity group is analyzed for insights to the analytic problem.
6. **Redefinition:** Activity groups need to be checked for consistency from time-to-time and may require redefinition based on new information.

Finally, while the Diamond Model is an analytic concept, the analytic outcomes are ultimately of value when applied to protecting assets and developing courses of action and mitigation strategy planning.

The model, due to its fundamental nature and inherent flexibility is useable by almost any mitigation planning and decision framework. For instance, it supports several critical steps of the Joint Intelligence Preparation of the Operational Environment (JIOPE) as well as Kill Chain analysis. The activity-attack graph already shows applicability to traditional information assurance planning models and the model has promising applications to even more advanced concepts such as the use of game theory and the development of strategies through evolutionary computing.

A Summary of Diamond Model Benefits

- Enables contextual and relationship-rich indicators improving cyber threat intelligence sharing and increasing the range of applicability of indicators
- Integrates information assurance and cyber threat intelligence through activity-attack graphs
- Improves analytic efficiency and effectiveness through easier identification of pivot opportunities and a simple conceptual method to generate new analytic questions
- Enhances analytic accuracy by enabling hypothesis generation, documentation, and testing, thereby applying more rigor to the analytic process
- Supports course of action development, planning/gaming, and mitigation strategies by integrating easily with almost any planning framework
- Strengthens cyber analysis tradecraft development by formalizing first principles upon which new concepts can be explored
- Identifies intelligence gap through a phase-based approach and the inclusion of external resource requirements as a fundamental meta-feature
- Supports real-time event characterization by mapping the analytic process to well-understood classification and intrusion detection research
- Establishes the basis of cyber activity ontologies, taxonomies, cyber threat intelligence sharing protocols, and knowledge management