The End of Semantic Diffusion: Why Cybersecurity Needs Its First Paradigm

Abstract

The cybersecurity industry exhibits characteristics of what Kuhn termed a pre-paradigmatic discipline: fragmented terminology, incompatible classification systems, and an inability to aggregate knowledge across institutions. This paper argues that cyber threats are causal and reproducible phenomena amenable to scientific treatment, yet the field lacks the shared taxonomic foundation required for progress. Drawing on Kuhn, Popper, Carnap, and Quine, I examine why semantic diffusion persists and articulate criteria for a unifying paradigm. I then propose a concrete instantiation: a cause-oriented taxonomy of ten non-overlapping threat clusters defined by the generic vulnerabilities they exploit, offering cybersecurity its first candidate paradigm.

1. Introduction

The world of cybersecurity presents a curious paradox: despite unprecedented investment in tools, expertise, and organizational capacity, defenders continue to lose ground against attackers. The conventional explanation—that adversaries are simply improving faster—obscures a deeper structural problem. This paper argues that cybersecurity suffers from a foundational deficiency that prevents cumulative progress: the absence of a shared paradigm.

Thomas Kuhn's *The Structure of Scientific Revolutions* (1962) provides a framework for understanding this condition. Kuhn described a phase in the development of scientific disciplines that precedes the establishment of shared foundations—what he termed the pre-paradigmatic phase. In this state, competing schools develop independent terminologies, methods, and explanatory models. They observe the same phenomena but cannot compare observations, combine findings, or build systematically upon each other's work. The breakthrough that transforms such a field into a mature science comes not through better instruments or more data, but through the establishment of shared axioms, definitions, and language.

I argue that cybersecurity currently occupies this pre-paradigmatic phase. The field exhibits precisely the symptoms Kuhn identified: fragmented terminology across major institutions, semantic diffusion in fundamental concepts, and an inability to translate findings into cumulative knowledge. Yet cyber threats possess characteristics that make them amenable to scientific treatment—they are causal, reproducible, and predictable. The question is not whether cybersecurity can become a science, but why it has not yet developed the shared foundation necessary to do so.

This paper proceeds as follows. Section 2 examines the symptoms of semantic diffusion in contemporary cybersecurity practice. Section 3 establishes that cyber threats follow scientific principles and are therefore suitable subjects for systematic inquiry. Section 4 draws on Popper, Carnap, and Quine to illuminate the epistemological requirements for disciplinary coherence. Section 5 articulates criteria for a unifying paradigm and proposes a

concrete instantiation: a cause-oriented taxonomy of ten threat clusters. Section 6 concludes with implications for the field's development.

2. The Symptoms of Semantic Diffusion

Contemporary cybersecurity exhibits four distinct symptoms of pre-paradigmatic disorder: fragmented terminology, semantic diffusion in core concepts, control-first reasoning, and the measurement problem.

2.1 Fragmented Terminology

Major cybersecurity organizations have developed independent classification systems that resist translation. CrowdStrike organizes threats around "eCrime" and "targeted intrusion" categories. Mandiant structures analysis by threat actor groups. The European Union Agency for Cybersecurity (ENISA) employs a distinct taxonomy. Verizon's Data Breach Investigations Report uses incident patterns as its organizing principle. Each organization observes the same attacks on the same digital infrastructure, yet their reports cannot be systematically compared, merged, or translated into consistent protective measures.

This fragmentation is not merely inconvenient; it is epistemologically debilitating. When CrowdStrike reports a 30% increase in "hands-on-keyboard intrusions" and Mandiant reports a 40% increase in "advanced persistent threats," we cannot determine whether they are describing the same phenomenon, overlapping phenomena, or distinct phenomena. The

absence of shared terminology prevents the aggregation of observations across institutions—precisely the condition Kuhn identified as characteristic of pre-paradigmatic disciplines.

2.2 Semantic Diffusion in Core Concepts

Beyond terminological fragmentation, cybersecurity suffers from semantic diffusion in its most fundamental concepts. The term "threat" is variously used to denote threat actors, vulnerabilities, techniques, control failures, incidents, and outcomes. "Risk" encompasses probability assessments, impact calculations, and general uncertainty. "Protection requirements" means different things across regulatory frameworks, organizational contexts, and technical domains.

When experts disagree about whether a particular event constitutes a "threat" or a "vulnerability" or a "control failure," they are not disagreeing about facts—they are operating with incompatible conceptual frameworks. This semantic diffusion makes reasoned disagreement impossible because practitioners cannot determine whether they are making contradictory claims about the same thing or compatible claims about different things.

2.3 Control-First Reasoning

A third symptom manifests in the structure of cybersecurity standards and regulations. Frameworks prescribe controls—technical safeguards, administrative procedures, physical protections—without systematically specifying the threats those controls are meant to address. This is analogous to a medical practice that prescribes treatments without diagnoses:

it may succeed occasionally, but it cannot learn from its successes or failures because it lacks the causal understanding that would explain either.

The consequences of control-first reasoning extend beyond practical inefficiency. Without explicit threat identification, organizations cannot assess whether their controls are appropriately targeted, whether observed failures indicate inadequate implementation or inappropriate control selection, or whether new threats require new control categories. The field accumulates compliance requirements without accumulating knowledge.

2.4 The Measurement Problem

The measurement problem in cybersecurity is a direct consequence of semantic diffusion. What cannot be named consistently cannot be measured reliably. When organizations use different definitions of "breach," "incident," "attack," and "compromise," aggregate statistics become meaningless. Each organization fights alone because there is no common basis for sharing experience.

This measurement problem distinguishes cybersecurity from mature security disciplines. Insurance actuaries can pool data across organizations because they share definitions of insurable events. Epidemiologists can track disease spread because they share case definitions. Cybersecurity cannot aggregate its experience because it lacks the shared conceptual foundation that would make aggregation meaningful.

2.5 The Failure of Canonical Frameworks

These symptoms pervade even the most authoritative frameworks in the field. The NIST Cybersecurity Framework (CSF 2.0), despite the word "cybersecurity" in its title, provides no cyber-specific threat definition and no threat categorization. It requires organizations to "identify threats" yet offers no taxonomy of what threats exist. The framework's threat categories in supporting documents (SP 800-30) are generic—Adversarial, Accidental, Structural, Environmental—categories that apply equally to physical security, fire safety, or any other risk domain.

ISO/IEC 27001, the most widely adopted information security standard, exhibits the same gap. Its 93 controls are organized by control type, not by threat addressed. Organizations implement controls without systematic guidance on which threats each control mitigates. MITRE ATT&CK, perhaps the most detailed threat knowledge base available, catalogs adversary techniques at the tactical level but provides no strategic categorization that would enable executive communication or cross-technique analysis. Its behavioral focus—what adversaries *do*—does not address the generic vulnerabilities they exploit.

The regulatory landscape compounds the problem. An analysis of thirty major cybersecurity regulations—NIS2, DORA, GDPR Article 32, HIPAA, PCI-DSS, and others across US, EU, UK, and Swiss jurisdictions—reveals a universal pattern: every regulation mandates controls without specifying the threats those controls address. This represents what might be called the logical impossibility of control-first regulation: frameworks that claim to

follow risk-based standards (which require threat identification before control selection) while systematically omitting the threat identification step.

The universality of this failure is itself significant. If even the canonical frameworks produced by NIST, ISO, and major regulatory bodies exhibit semantic diffusion, the problem is not institutional negligence but disciplinary immaturity. The field lacks the shared taxonomic foundation that would make threat-specific control mapping possible.

3. Cyber Threats as Scientific Phenomena

One might object that cybersecurity is not a "real" science—not a natural science—and therefore cannot be expected to develop the kind of shared foundation that characterizes physics or chemistry. This objection misunderstands what makes a domain amenable to scientific treatment.

Cyber threats exhibit two characteristics that make them suitable subjects for scientific inquiry: causality and reproducibility. A buffer overflow follows the same principles whether exploited in Berlin or Bangalore. SQL injection operates according to identical mechanisms regardless of who executes it. Social engineering exploits universal cognitive patterns rooted in human psychology. These are not random events or one-time occurrences; they are systematic exploitations of regularities.

Where causality reigns and reproducibility is possible, where predictions can be made and tested, the methods of scientific inquiry apply. The question is not whether cybersecurity

follows scientific principles—it manifestly does. The question is why the field has not yet established the shared foundation that would enable scientific progress.

The history of science provides instructive parallels. Physics before Newton, chemistry before Lavoisier, medicine before germ theory—each possessed brilliant practitioners, useful techniques, and partial knowledge. What each lacked was the taxonomic foundation that transforms craft into science. Cybersecurity appears to occupy an analogous position: substantial practical capability without the conceptual infrastructure for cumulative progress.

4. Epistemological Requirements for Disciplinary Coherence

To understand what cybersecurity requires to achieve disciplinary coherence, it is useful to draw on three philosophers of science whose work addresses problems the field encounters daily, even if practitioners rarely frame them in these terms.

4.1 Popper and Falsifiability

Karl Popper's (1959) demarcation criterion holds that a theory is scientific only if it is falsifiable—only if it makes claims that could, in principle, be refuted by observation. A statement compatible with every possible observation tells us nothing about the world. This criterion has implications for how cybersecurity should structure its claims.

A coherent threat taxonomy must make testable claims. If it asserts that all cyber attacks can be classified into a finite set of categories, then the discovery of an attack that cannot be so classified would refute it. This is not a weakness but a strength: the possibility of refutation is what distinguishes genuine knowledge from mere speculation. A taxonomy that claims to classify "threats" but defines the term so broadly that any event might count as a threat is not falsifiable and therefore not scientifically useful.

4.2 Carnap and Logical Syntax

Rudolf Carnap (1934), a central figure in the Vienna Circle, devoted his career to the question of how to construct language in which misunderstandings become structurally impossible. His answer involved formal systems—syntactically precise, semantically unambiguous.

Carnap would likely diagnose cybersecurity's semantic diffusion as a symptom of insufficiently formalized language. When "threat" means ten different things in ten different contexts, the problem is not expertise but grammar—the field lacks a shared syntax that would constrain interpretation. The remedy is not better training but better definitions: foundational terms must be specified so precisely that they remain stable across contexts. This is not pedantry; precise communication is the precondition for collective action.

4.3 Quine and Holism

W.V.O. Quine (1951) challenged the notion that individual statements can be verified in isolation. In "Two Dogmas of Empiricism," he argued that our beliefs form an interconnected web. When an experiment fails, we can adjust various nodes in the web—not merely the hypothesis most directly tested.

For cybersecurity, this holism is an uncomfortable truth. When an attack succeeds, multiple factors typically contribute: firewall configuration, process design, human decision-making, architectural choices. Responsibility distributes across the web of organizational practice. A coherent threat framework does not deny this complexity but provides structure for naming the nodes. Only when we can identify which factors interact can we meaningfully analyze their interrelationships and design appropriate interventions.

5. A Proposed Paradigm: Cause-Oriented Threat Clusters

Drawing on the preceding analysis, I now propose a concrete paradigm for cybersecurity: a cause-oriented taxonomy that classifies threats by the generic vulnerabilities they exploit rather than by the outcomes they produce. This taxonomy—which I term Top Level Cyber Threat Clusters (TLCTC)—is designed to satisfy the epistemological requirements articulated above while providing practitioners with an immediately applicable framework.

5.1 The Foundational Principle: Causality Over Outcomes

The fundamental error of existing taxonomies is that they classify by outcomes—ransomware, data breach, denial of service—rather than by causes. This is analogous to a medical taxonomy that groups diseases by symptoms (fever, pain, fatigue) rather than by etiology (bacterial infection, viral infection, autoimmune disorder). Symptom-based classification conflates distinct causal mechanisms and obscures the relationship between diagnosis and treatment.

I propose that cyber threats should be classified by the *generic vulnerability* they exploit—the root weakness that makes the attack possible. Generic vulnerabilities are universal: they persist across system types, software implementations, and evolving attack techniques. A buffer overflow exploits the same generic vulnerability (implementation flaws in server-side code) whether it targets a web application in 2005 or a cloud service in 2025. This stability is what makes cause-oriented classification scientifically tractable.

5.2 The Ten Threat Clusters

I propose that all cyber threats can be classified into exactly ten non-overlapping clusters, each defined by a distinct generic vulnerability. This claim is falsifiable: if an attack emerges that exploits a generic vulnerability outside these ten categories, the taxonomy must be revised. The clusters are:

Cluster 1: Abuse of Functions. Generic vulnerability: the scope of software functions and features. Attacks in this cluster do not exploit implementation flaws; they misuse legitimate functionality for unintended purposes. Examples include business email compromise, configuration manipulation, and abuse of administrative tools. The system works as designed—the design merely permits harmful use.

Cluster 2: Exploiting Server. Generic vulnerability: implementation flaws in server-side code. This cluster encompasses SQL injection, buffer overflows, authentication bypasses, and other attacks that exploit coding errors in systems that receive and process requests. The attack vector is always directed at the server role in a client-server interaction.

Cluster 3: Exploiting Client. Generic vulnerability: implementation flaws in client-side code. This cluster covers attacks that exploit vulnerabilities in software occupying the client role—browsers, document readers, email clients, media players. A malicious PDF that exploits a parsing flaw in the reader is classified as Cluster 3; the code execution that follows belongs to Cluster 7, yielding the chain #3→#7. This distinction matters: Cluster 3 is addressed by patching client software and sandboxing parsers, while Cluster 7 is addressed by execution controls and endpoint protection. Collapsing both into "malicious document" obscures these distinct control requirements.

Cluster 4: Identity Theft. Generic vulnerability: weaknesses in identity and credential management. A critical distinction governs this cluster: credential *acquisition* versus credential *use*. Acquiring credentials through phishing is classified #9→#4 (social engineering enables identity compromise). Acquiring credentials through a server exploit is #2→#4. Acquiring credentials through malware is #7→#4. However, *using* already-obtained credentials—credential stuffing, pass-the-hash, token replay—begins at Cluster 4 as the initial vector. The generic vulnerability (reliance on identity artifacts that can be replayed) is exploited directly. This distinction enables precise control mapping: preventing acquisition requires controls at #9, #2, or #7; preventing misuse requires controls at #4 (multi-factor authentication, session binding, anomaly detection).

Cluster 5: Man in the Middle. Generic vulnerability: lack of end-to-end communication protection. This cluster covers attacks that exploit a controlled position on a communication path—traffic interception, protocol downgrade attacks, and active manipulation of data in transit. The attack requires the adversary to have already obtained a position enabling interception; *gaining* that position is classified separately (often #1 for network reconfiguration or #8 for physical access to network infrastructure).

Cluster 6: Flooding Attack. Generic vulnerability: finite capacity limitations inherent in any system component—bandwidth, CPU, memory, storage, quotas, connection pools. This cluster covers exhaustion of resources through volume or intensity that exceeds

capacity limits. Three boundary tests govern classification: (1) If availability loss results primarily from an *implementation defect*—a crash, an algorithmic complexity weakness such as ReDoS—the attack belongs to Cluster 2 or 3, not Cluster 6, because the exploited vulnerability is a code flaw, not finite capacity. (2) If availability loss results from *capacity exhaustion by volume or intensity*—SYN floods, amplification attacks, volumetric DDoS—the attack is Cluster 6. (3) If attackers *amplify load by abusing legitimate functions*, the enabling step may be Cluster 1, but the exhaustion event remains Cluster 6, yielding the chain #1→#6. These boundary tests illustrate how cause-oriented classification resolves ambiguities that outcome-based taxonomies cannot: three scenarios produce identical unavailability but require fundamentally different controls.

Cluster 7: Malware. Generic vulnerability: designed code execution capabilities. Systems are built to execute programs; malicious programs exploit this design. Importantly, Cluster 7 as an *initial* vector is relatively rare in sophisticated attacks. More commonly, code execution *follows* another cluster: #3→#7 (exploit enables execution), #9→#7 (user tricked into running malware), #10→#7 (supply chain delivers malicious code). Cluster 7 is the initial vector when execution capabilities are targeted directly—autorun from removable media, exploitation of automatic execution settings, or environments where code execution requires no prior compromise. This analysis reveals why "malware" as a threat category is analytically inadequate: it names the payload but not the delivery vector, conflating attacks with fundamentally different prevention requirements.

Cluster 8: Physical Attack. Generic vulnerability: physical accessibility of systems. Hardware implants, evil maid attacks, USB drops, and physical theft exploit the fact that digital systems exist in physical space. Physical access enables attack chains: #8→#7 (removable media introduces malware), #8→#5 (physical access to network infrastructure enables interception). The cluster marks the boundary crossing from physical domain to cyber domain.

Cluster 9: Social Engineering. Generic vulnerability: human psychological factors. Phishing, pretexting, and manipulation attacks exploit cognitive biases and trust relationships universal to human psychology. Cluster 9 is typically an *enabler* that initiates chains to other clusters: #9→#4 (manipulation yields credentials), #9→#7 (user executes malicious attachment), #9→#1 (employee tricked into unauthorized wire transfer). The chain notation reveals what "phishing" alone obscures: the same social engineering technique can enable completely different downstream attacks with completely different consequences. Classifying all phishing identically—as outcome-based taxonomies do—conflates attacks requiring different defensive responses.

Cluster 10: Supply Chain Attack. Generic vulnerability: third-party trust dependencies. Organizations necessarily depend on software, hardware, and services they do not fully control. This cluster applies when the *victim's initial exposure* comes through a compromised trusted source. Consider the SolarWinds incident: from the perspective of

affected organizations, their initial vector was #10—they received malicious code through a trusted software update channel. The attacker's chain *against SolarWinds itself* was different (likely involving #9, #4, #2, or #7 to compromise the build environment). This perspective-dependence is a feature, not a bug: each organization's risk analysis must begin with *its own* initial exposure. Victims of supply chain attacks face Cluster 10; preventing such attacks requires controls at the supplier's exposure points.

5.3 Axioms and Classification Rules

To ensure consistent application, the taxonomy operates under explicit axioms that constrain interpretation:

Axiom I: No System-Type Differentiation. The taxonomy applies uniformly across IT system types. Whether the target is enterprise IT, cloud infrastructure, operational technology, IoT devices, or mobile endpoints, the same ten generic vulnerabilities apply. Sector labels do not create new threat classes; they merely specify which concrete vulnerabilities instantiate the generic categories in particular contexts.

Axiom II: Client-Server as Universal Model. Any networked system interaction can be modeled as a client-server relationship at some level of abstraction. This axiom grounds the distinction between Clusters 2 and 3: the same implementation flaw may be classified differently depending on whether the vulnerable code occupies the server or client role in the interaction.

Axiom III: Threats Are Causes, Not Outcomes. Threat clusters occupy the causal side of risk analysis. Outcomes—loss of confidentiality, integrity, or availability—are recorded separately as consequences. "Data breach" is not a threat; it is a possible outcome of threats from multiple clusters.

Axiom IV: Threats Are Not Threat Actors. The taxonomy classifies actions and vulnerabilities, not actors. Attribution, motivation, and capability are analytically separate from the classification of how an attack operates. The same threat cluster applies whether the attacker is a nation-state, a criminal organization, or an insider.

Axiom V: Control Failure Is Not a Threat. The absence or failure of a control is a condition that may enable a threat, but it is not itself a threat category. This axiom prevents the conflation of vulnerability assessment with threat classification.

5.4 Attack Paths and Sequential Analysis

Real-world attacks rarely exploit a single vulnerability in isolation. The taxonomy accommodates multi-stage attacks through the concept of *attack paths*—ordered sequences of cluster steps that trace an intrusion from initial access to final impact. Each step in the path exploits exactly one generic vulnerability and therefore maps to exactly one cluster.

Consider a sophisticated intrusion: an attacker sends a phishing email (Cluster 9) containing a malicious document that exploits a rendering vulnerability (Cluster 3), establishing a foothold from which malware executes (Cluster 7), enabling credential theft (Cluster 4), which permits lateral movement through legitimate administrative tools (Cluster

1), ultimately resulting in data exfiltration. This attack path—9→3→7→4→1—is more analytically precise than labels like "advanced persistent threat" or "targeted intrusion" because it specifies which generic vulnerabilities were exploited in what sequence.

Attack path notation enables comparative analysis across incidents. Organizations can identify which clusters appear most frequently in their incident data, which transitions between clusters are most common, and which points in typical attack paths offer the most leverage for defensive intervention.

5.5 Separation of System Risk and Data Risk Events

The taxonomy introduces a critical distinction that existing frameworks typically conflate: the difference between *system risk events* (loss of control over a system) and *data risk events* (consequences for data confidentiality, integrity, or availability). In the Bow-Tie model of risk analysis, threat clusters occupy the left (causal) side; data risk events occupy the right (consequence) side. The central event—system compromise—is the pivot between them.

This separation has practical implications. Controls on the causal side aim to *prevent* system compromise; controls on the consequence side aim to *mitigate* data risk events after compromise has occurred. Conflating these—treating "data breach" as both a threat and an outcome—obscures the distinction between prevention and mitigation and makes control mapping incoherent.

5.6 Satisfaction of Epistemological Criteria

The proposed taxonomy satisfies the epistemological requirements derived from Popper, Carnap, and Quine:

Falsifiability. The claim that all cyber threats exploit one of ten generic vulnerabilities is testable. Discovery of an attack that exploits a genuinely novel generic vulnerability—one not reducible to the existing categories—would refute the taxonomy or require its revision. This susceptibility to refutation distinguishes the proposal from unfalsifiable frameworks that can accommodate any observation through post hoc reinterpretation.

Semantic Precision. The explicit axioms and classification rules function as Carnap's logical syntax—formal constraints that reduce interpretive ambiguity. "Threat" is defined as a causal category, distinct from actors, outcomes, and control failures. Each cluster has a specified generic vulnerability that determines membership. Independent practitioners applying these rules to the same incident should arrive at the same classification.

Holistic Integration. Attack path notation addresses Quine's insight that understanding distributes across interconnected elements. Rather than treating each attack step in isolation, the taxonomy provides structure for analyzing how threats combine—which cluster transitions are common, which combinations are particularly dangerous, which defensive interventions disrupt the most attack paths.

Comparability. When multiple organizations adopt the same taxonomy, their observations become aggregable. An attack classified as 9→3→7→4→1 by one organization can be compared directly with similarly classified attacks from other organizations. This comparability is precisely what fragmented terminology currently prevents.

6. Conclusion

Kuhn observed that paradigms do not establish themselves by decree; they establish themselves by solving problems that previously seemed intractable, by enabling communication that was previously impossible, by providing practitioners with tools that demonstrably work. The anomalies of cybersecurity's current state are manifest: reports that cannot be compared, standards that prescribe controls without naming threats, experts who talk past each other for lack of shared language.

The taxonomy proposed here—ten cause-oriented threat clusters, governed by explicit axioms, supporting sequential attack path analysis, and distinguishing system risk from data risk—offers cybersecurity a candidate for its first paradigm. It makes falsifiable claims, provides semantic precision, accommodates holistic analysis, and enables cross-organizational comparison.

Whether this particular proposal achieves adoption is less important than whether the field recognizes the need for *some* shared paradigm. The transition from pre-paradigmatic to paradigmatic status requires what Kuhn called a "conversion experience"—a shift in how practitioners see their domain. Such shifts cannot be forced, but they can be prepared for.

History suggests that disciplines resist paradigm shifts until the inadequacy of existing frameworks becomes undeniable. Cybersecurity may be approaching that threshold.

The proposal is offered in the spirit of Popper's falsificationism: use it, test it, attempt to break it—and report what fails. If the ten clusters prove insufficient, if attacks emerge that cannot be classified, if the axioms generate contradictions, these discoveries will advance understanding more than adherence to a flawed framework ever could. Science progresses through refutation; cybersecurity can do the same, once it has something precise enough to refute.

References

Carnap, Rudolf. 1934. *Logische Syntax der Sprache*. Vienna: Springer. Translated as *The Logical Syntax of Language* (London: Kegan Paul, 1937).

Kuhn, Thomas S. 1962. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.

Popper, Karl R. 1959. *The Logic of Scientific Discovery*. London: Routledge.

Quine, W.V.O. 1951. "Two Dogmas of Empiricism." *Philosophical Review* 60:20–43. https://doi.org/10.2307/2181906