PRE-PUBLICATION DRAFT. This is the author's accepted version of an article accepted for publication in *IEEE Security & Privacy*. It has not been copyedited or typeset by IEEE, and the published version will differ in wording and layout. Please cite the published version: E. H. Spafford, "What Do We Mean by 'Cybersecurity'? Definitional Ambiguity and Its Consequences for Measurement, Policy, and Progress," *IEEE Security & Privacy*, to appear, Jan./Feb. 2027 (special issue on Cyber Hard Problems), doi: 10.1109/MSEC.2026.3748675.

© 2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

What Do We Mean by "Cybersecurity"?

Definitional Ambiguity and Its Consequences for Measurement, Policy, and Progress

Eugene H. Spafford, Life Fellow, IEEE

E. H. Spafford is with the Department of Computer Science and the Center for Education and Research in Information Assurance and Security (CERIAS), Purdue University, West Lafayette, IN 47907 USA (email: spaf@purdue.edu).

Manuscript prepared April 2026; revised September 2026.

Abstract — After decades of practice, cybersecurity still lacks a consensus definition, which hampers measurement, regulation, and insurance. Drawing on how aviation, nuclear safety, and quality management converged, this article proposes and tests a functional definition and argues for a common core with context-specific operationalization.

Index Terms — cybersecurity, security metrics, risk management, security policy, cybersecurity definition, trust, privacy, cyber insurance, workforce

Introduction: Everyone Knows What "Cybersecurity" Means

In an era when cybersecurity commands headlines, billions of dollars in annual spending, and the attention of heads of state, one might reasonably assume that the field rests on a firm definitional footing. It does not. After nearly seven decades of study and practice, no commonly accepted, precise definition of cybersecurity exists.

In the period 1960–2000, practitioners spoke of "computer security," "communications security," "information security," "network security," and "data security." Those terms were compact and descriptive, but discussing an integrated view required awkward compounds such as "computer and network and data security," which is unwieldy and resists acronymic compression. In the late 1990s, when a U.S. Senate committee held hearings on the security of government systems, a staffer apparently coined "cybersecurity" as shorthand. Perhaps it seemed suitably exotic to the Senators. In any case, it caught on, much to the dismay of many working professionals.[a]

However, "cyber" is imprecise. It is easy to lose sight of the data, processes, people, organizations, and policies that the term is meant to encompass, in addition to computers and networks.

That vagueness shows up in formal definitions. The online Merriam-Webster dictionary defines cybersecurity as "measures taken to protect a computer or computer system (as on the Internet) against unauthorized access or attack," omitting the control of *authorized* access and the protection of networks. The U.S. National Institute of Standards and Technology (NIST) defines cybersecurity far more broadly:

[a] I had a conversation about this with Senator Sam Nunn at the time, expressing reservations about the term. As the Senator's chief of staff replied, "It's you versus a senior U.S. Senator on this. You lose."

"Prevention of damage to, protection of, and restoration of computers, electronic communications systems, electronic communications services, wire communication, and electronic communication, including information contained therein, to ensure its availability, integrity, authentication, confidentiality, and nonrepudiation." [1] This is also the definition adopted by U.S. Department of Defense Instruction 8500.01. However, NIST maintains at least six distinct definitions of "cybersecurity" in its Computer Security Resource Center glossary, derived from different authoritative sources and contexts [1]. NIST has not consolidated them.

The Internet Engineering Task Force's RFC 4949 offers yet another formulation: "A system condition in which system resources are free from unauthorized access and unauthorized or accidental change, destruction, or loss." Rob Joyce offered a shorter definition in an NSA newsletter in 2019: "Cybersecurity is everything that results in protecting information and underlying technology from theft, manipulation, and disruption." An even more succinct formulation appeared decades earlier, in 1991, from Garfinkel and Spafford: "A computer is secure if you can depend on it and its software to behave as expected." [2]

International bodies add to the proliferation. The EU Cybersecurity Act (Regulation 2019/881) defines cybersecurity as "the activities necessary to protect network and information systems, the users of such systems, and other persons affected by cyber threats," [3] broadening the scope to include affected individuals, not merely systems. The NIS2 Directive (2022/2555) references this definition rather than formulating its own [4]. ISO/IEC 27032:2023 declines to provide a standalone definition, instead deferring to related terms in the ISO/IEC 27000 family. The 2023 U.S. National Cybersecurity Strategy characterizes cybersecurity as essential to the economy, critical infrastructure, and democracy, without providing a formal definition.

Craigen, Diakun-Thibault, and Purse observed in their systematic analysis that the term "cybersecurity" exhibits "highly variable, context-bound, often subjective, and, at times, uninformative" definitions [5].

These differences matter because without an agreed-upon definition, everything else is imprecise: our metrics, our comparisons, our policies, and our claims of progress. A shared definition cannot make threat environments, risk appetites, or legal regimes alike, nor will it repair a badly designed metric. What it can do is make claims about those things comparable.

Part of the confusion is categorical. Some of these definitions treat cybersecurity as a *property* (a state or condition of a system), while others treat it as a *discipline* (the practice of achieving and maintaining that state). This article uses the term in both senses, as context requires. The definition proposed later clarifies their relationship.

What a Definition Can and Cannot Fix

Many cybersecurity professionals cringe at the question "How secure are we?" The person asking probably expects a number, perhaps a percentage or a rating. No such answer is meaningful. Without a clear definition, we cannot form proper metrics.

Textbooks and courses often invoke the traditional C-I-A triad (confidentiality, integrity, and availability) as the basic components of cybersecurity. But these are poorly suited as metrics. What is the dimension of integrity? How does one have "two more units of confidentiality"? Are three units of confidentiality more important than two units of availability? These properties are not orthogonal, either. If data is overwritten (a failure of integrity), it is no longer available. The shortcomings of the C-I-A model are not new realizations. Subsequent models (Parker's Hexad, McCumber's Cube, the Clark-

Wilson integrity model, Microsoft's STRIDE threat-modeling framework, and later extensions) attempted to address these shortcomings but did not resolve the underlying definitional problem.

These models are valuable, but they operate atop a definitional void. Three distinct problems are at work, and they require different remedies. First, there is a *metric-design* problem. The C-I-A triad and its successors are technically flawed as measurement instruments, regardless of how "cybersecurity" is defined. Second, there is a *comparability* problem. Because each model embeds its own working definition of what it means to be "secure," results from one are not directly comparable to results from another. Third, there is a *conceptual* problem. Without agreement on what cybersecurity *is*, it is unclear what we should be measuring and teaching. The second and third are definitional problems, and a shared definition would address them. The first is an engineering problem that no definition will solve directly. A fourth characteristic, *operational variation*, is not a problem at all. A bank and a hospital face different threats, tolerate different risks, and answer to different regulators. Their measurements should therefore differ. However, all four issues are routinely conflated, which obscures where a definition would help and where it would not.

Two facts not reflected in most practice complicate all of this. First, *it is impossible to make an arbitrary system secure against all threats and still keep it usable, functional, and productive*. Or, as in my often-cited quote from 1989, "The only truly secure system is one that is powered off, cast in a block of concrete and sealed in a lead-lined room with armed guards—and even then I have my doubts." Imagine protecting a PC not only against the combined cyberattack capabilities of NATO forces and every cybercriminal gang, but also against a planet-killing asteroid. Those are potential threats, and there is no current defense against all of them while maintaining fully useful functionality. Second, *all security is relative to a security policy*. A graduate student's recipe for chocolate chip cookies and a multinational bank's balance sheets may reside on identical hardware running identical operating systems, but the risks, policies, and appropriate controls differ. One size does not fit all.

Consequences of Definitional Failure

Definitional ambiguity is not merely a philosophical nuisance. It has concrete consequences for law, regulation, insurance, and workforce development.

Legal, Regulatory, and Insurance Fragmentation

When different jurisdictions define cybersecurity differently, the resulting regulatory fragmentation burdens organizations and complicates enforcement.

In the United States, all 50 states plus the District of Columbia and U.S. territories have enacted security breach notification laws. These laws do not use consistent definitions of "breach," "personal information," or "cybersecurity incident." Harmonizing compliance among these state-level regimes is nearly impossible without preemptive federal legislation, and every Congressional attempt at a federal breach-notification standard has stalled.

The SEC's 2023 cybersecurity disclosure rules define a "cybersecurity incident" as an unauthorized occurrence that jeopardizes the confidentiality, integrity, or availability of a registrant's information systems. The word "unauthorized" does interesting work here. It potentially excludes incidents caused by negligence, ineptitude, misconfiguration, or system failure, events that many practitioners would consider security-relevant. The EU's NIS2 Directive [4] defines a "significant incident" by reference to operational disruption and financial loss, a different axis entirely. The Financial Stability Board (FSB) confronted this

divergence directly in its 2023 update to its Cyber Lexicon, explicitly adopting a broad definition of "cyber incident" that includes non-malicious events "to promote convergence." That a global financial regulatory body had to take an explicit position on what "incident" means should give us pause.

Here the definitional and operational strands separate cleanly. A shared definition of "incident" would tell the SEC and the EU whether they are counting the same events. It would not, and should not, impose the same reporting threshold on a community bank and a stock exchange. "Significant" is a threshold, and sectors and jurisdictions properly set their own. The comparability problem is definitional. The threshold differences are operational variation, and they would remain after convergence but would be more easily reconciled.

The OECD's 2024 report on measuring cybersecurity concluded that the absence of agreed definitions makes cross-national comparison of cybersecurity posture effectively impossible [6].

The cyber insurance market also illustrates the problem. Insurers must underwrite a risk they cannot consistently define. A 2021 GAO report stated that "developing cyber insurance products is difficult because insurers don't have much historical data on cyberattack-related costs, and determining what's covered is hard for clients because key terms like 'cyberterrorism' don't have standard definitions." [7] The GAO's two findings differ. The missing loss history is an actuarial problem that only time and data will cure, whereas the missing definitions could be addressed with a standard definition.

The Workforce We Cannot Count

If we cannot define cybersecurity, we cannot define a cybersecurity professional or prescribe what they should be taught. If we cannot define professionals, we cannot count them, which means the widely cited "workforce gap" deserves scrutiny.

The NICE Workforce Framework (NIST SP 800-181 Rev. 1) attempts to decompose cybersecurity work into tasks, knowledge areas, and skills organized into seven work-role categories. It is a useful tool, but it necessarily embeds its own working definition of what counts as "cybersecurity work."

ISC2's 2024 Cybersecurity Workforce Study estimated a global workforce gap of 4.76 million people, with 47% of the estimated need unmet. However, that estimated "need" depends entirely on how one defines cybersecurity work. A broad definition produces a large gap, while a narrow one produces a smaller gap. ISC2 itself appeared to acknowledge this. Its 2025 report shifted away from gap estimates, asserting that skills shortages matter more than headcount shortages. That is progress, but it sidesteps the underlying definitional question. Most of the definitions surveyed above name computers, networks, and information, but few name the people who operate them or the organizations that depend on them. If our definitions undercount the human aspects of cybersecurity, our workforce models will too.

This is the conceptual problem in its purest form. The count depends on the definition. A shared definition would not tell any employer how many people to hire, and it would not erase the differences between what a bank and a hospital need. Those are operational matters. It would make workforce estimates commensurable, and it would say what the people being counted are supposed to know.

Other Fields Converged, but Cybersecurity Differs

Other fields offer useful comparisons, provided their differences from ours are stated first. Aviation safety, nuclear safety, and quality management have stable physical referents such as aircraft, reactors, and manufactured goods. Their hazards do not strategize against defenders. Cybersecurity has neither a stable referent nor hazards as constant as gravity and material fatigue. Its adversaries are intelligent,

motivated, and quick to adapt to whatever defenses are deployed. Thus, the threat environment shifts in response to defensive choices, and any definition risks going stale. These are differences in kind, not merely in degree, and some of the field's definitional fragmentation is a rational response. The comparisons that follow therefore claim less than they might appear to. They show that convergence is achievable in complex, contested, safety-critical domains, not that cybersecurity can follow the same path at the same pace. The adversarial dimension is not unique, either. Aviation safety must account for terrorism, nuclear safety for sabotage and theft, and quality management for counterfeit parts and supply-chain fraud. What distinguishes cybersecurity is the *tempo* of adversarial adaptation, and that argues against a static definition, not against a shared one.

Aviation provides one clear example. In 2013, the International Civil Aviation Organization (ICAO) adopted Annex 19 to the Convention on International Civil Aviation, which defines safety as "the state in which risks associated with aviation activities… are reduced and controlled to an acceptable level." [8] That definition is now binding on all 193 ICAO member states. It is not a perfect definition (one can quibble with "acceptable level"), but it provides a shared foundation on which metrics, regulatory regimes, and international cooperation can be built.

Nuclear safety followed a similar trajectory. The International Atomic Energy Agency (IAEA) defines safety as "the protection of people and the environment against radiation risks, and the safety of facilities and activities that give rise to radiation risks." [9] The IAEA Safety Glossary, first published in 2000 and now in its 2022 interim edition, provides a shared vocabulary for 176 member states.[b] The definition is broad enough for different national contexts yet precise enough to anchor inspections, licensing, and international cooperation.

Even the notoriously abstract concept of "quality" has been formalized. ISO 9000:2015 defines quality as "the degree to which a set of inherent characteristics of an object fulfills requirements." [10] The path from the first ISO 9000 publication in 1987 to the structurally harmonized 2015 edition took 28 years.

These examples share a common structure. Each involved a single authoritative international body (ICAO, IAEA, ISO Technical Committee 176) that convened multidisciplinary expertise and had the institutional standing to drive adoption. Each definition is deliberately *functional*. It says what the discipline intends to achieve, not how it gets there. Each accepts that the concept it defines is relative ("acceptable level," "fulfills requirements") while providing sufficient precision to support measurement and comparison. Each leaves thresholds to be set locally. Each required a sustained, multi-year effort.

Cybersecurity has none of these structures. No single authoritative international body has the standing and mandate to drive definitional convergence. NIST, ISO, ENISA, the ITU, and dozens of national agencies all define cybersecurity independently, often inconsistently. The field is more multidisciplinary than aviation safety or quality management—computer science, law, psychology, management, political science, and philosophy all claim authority—which makes consensus more difficult. However, "more difficult" is not "impossible," and the examples above suggest that the effort could succeed. As long as we lack a shared definition, we will keep using metrics that cannot be compared, policies that cannot be harmonized, and progress that cannot be measured.

[b] IAEA membership as of the 2022 interim edition of the glossary.

What a Useful Definition Must Do

Functional: Security as Instrument, Not End

Some people, particularly cybersecurity professionals, believe others want cybersecurity for its own sake. This view may sound reasonable, but it is misguided. *Cybersecurity is not the primary goal but a means of maximizing and supporting what users are trying to accomplish.* People and organizations have goals that cybersecurity can protect and enable, but their primary tasks matter most. Users want to buy things online. Hospitals want to treat patients. Gamers want to play games. Cybersecurity supports these goals by protecting the user and the activity from adversity and loss.

This framing has direct implications for the definitional problem. A useful definition of cybersecurity must be *functional*. It must characterize cybersecurity by its role in enabling and protecting primary missions, not as a state to be achieved for its own sake.

When security displaces the primary mission, the cost can be measured, sometimes in lives. As one example, Choi, Johnson, and Lehmann [11] examined the relationship between data-breach remediation and hospital care quality. Because cybersecurity measures often increase overhead after a breach, the researchers asked whether additional security measures increased the time from a patient's arrival in the emergency department to receiving an electrocardiogram (EKG). Their data showed an increase of 0.5 to 2.7 additional minutes in the three years following a breach, with associated increases in 30-day mortality from acute myocardial infarction.

The takeaway is not to focus on maximizing *security*. We should focus on protecting the users' primary goals with an appropriate amount of security, calibrated to context and risk tolerance.

Context, Policy, and Trust

Two facts stated earlier—that absolute security is impossible and that all security is relative to policy—were well understood decades ago. Buried in RFC 4949 is a gem, *Courtney's Laws*, articulated by Robert H. Courtney, Jr., a pioneer in cybersecurity:

> *Courtney's first law:* You cannot say anything interesting (i.e., significant) about the security of a system except in the context of a particular application and environment.
>
> *Courtney's second law:* Never spend more money eliminating a security exposure than tolerating it will cost you.
>
> - First corollary: Perfect security has infinite cost.
> - Second corollary: There is no such thing as zero risk.
>
> *Courtney's third law:* There are no technical solutions to management problems, but there are management solutions to technical problems.

Courtney's first law is an operational constraint. A definition that cannot accommodate the vast differences between protecting a student's personal laptop and protecting critical national infrastructure is not useful. Context includes not only the assets and threats but also the organizational culture, available resources, legal environment, and risk tolerance. Cybersecurity practitioners sometimes assume that what matters to them matters to everyone else. It does not. Context is also where legitimate variation lives. A definition that accommodates context anticipates sectors and jurisdictions operationalizing it differently.

Those who have studied the field in depth generally prefer to speak of *trust* or *risk* rather than "security" in the abstract. One of the most influential documents in the field was the Trusted Computer

System Evaluation Criteria (TCSEC), colloquially "the Orange Book," issued by the U.S. National Computer Security Center in 1983. The choice of language was deliberate, because it makes more sense to ask how much we trust a system to operate in accordance with policy than to claim the system "is secure."

There is a vast difference between threat-based and risk-based security. "Our firewall prevents network attacks" is qualitatively different from "our firewall lowers the risk of network attacks." As Schneier observed, "Avoiding threats is black and white: either you avoid the threat, or you don't. Avoiding risk is continuous: there is some amount of risk you can accept, and some amount you can't." [12] Risk management is a spectrum from none to complete, with zero risk being unattainable.

Both point to the centrality of *security policy*, the document (or understanding) that defines what we are protecting, from whom, and at what cost. Few organizations have an overall security policy that can be directly translated into specifications for developing, procuring, and operating their computing resources. Most effectively operate under a policy of "no one should do anything we do not authorize." They depend on off-the-shelf components to enforce it, although those components were never carefully specified or designed to support the full range of possible policies.

As Young, Boebert, and Kain observed in 1985, a system without requirements cannot fail; it merely presents surprises [13]. Many security "bugs" and "weaknesses" are thus not flaws in the usual sense—they are *surprises*, occurring in systems that were never fully defined and specified against security policies. If we have no policy, and the software we run was never explicitly designed against one, how can there be a security "flaw"?

Trends involving just-in-time development, DevOps, and Agile methodologies often shortchange detailed design in favor of speed of delivery. The philosophy is that problems can be quickly fixed, but *ease or speed of patching is not the same as good security.*

A definition must be stated relative to context and policy, and it should speak of risk and trust rather than of security as an absolute.

Support for Measurement

A definition must support measurement. This does not mean that a definition must itself provide metrics, but it must be sufficiently precise to ground the development of metrics. The proliferation of frameworks (FAIR, NIST CSF, ISO 27001, CIS Controls), each embedding its own implicit definition, means an organization's "security posture" depends substantially on which framework it uses.

Despite lists of "top 10" best practices, there are no accepted values for precisely how much patching, firewalls, or training lower risk. We know these measures help "to an extent," but the remaining risk lies within the enormous range from 0% to 100%. A definition will not fix the metric-design problem, but it would clarify what the metrics are supposed to measure.

Coverage of Non-Obvious Risks

Cybersecurity is not always about the apparent risks or the technology itself. A definition that focuses exclusively on computers, networks, and software misses critical dimensions.

The dominant framing of cybersecurity risks leans heavily on external threats, epitomized by the hooded figure hunched over a keyboard. But many attacks originate from within the organization. Insiders do not need remote code execution as they are already local. A definition of cybersecurity that privileges external-threat defense over insider risk, organizational culture, or human factors is incomplete.

This connects to *fear, uncertainty, and doubt* (FUD) in cybersecurity. Vendors and pundits highlight threats so often that there is an acronym for it: FUD. Warnings have value, but fear-based approaches are counterproductive. As Ciarán Martin, the first CEO of the UK's National Cyber Security Centre, has noted, the field must move away from a fear-based approach toward a pragmatic one. A good definition should support pragmatic risk management, not fear-driven spending.

Technology Neutrality

Artificial intelligence, quantum computing, and whatever comes next should not fall outside any reasonable definition of cybersecurity, nor should their arrival break it. A definition so tightly coupled to today's threat environment that it becomes obsolete with the next major disruption is a snapshot, not a definition. The next section uses AI as the test of this requirement.

Ethics, Privacy, and Safety: Where They Belong

Three further concerns recur in discussions of what cybersecurity is: ethics, privacy, and safety. One objection is that these are principles that should guide practice rather than elements of a definition. Whether that holds differs for each.

Privacy belongs inside the definition, as one of the interests the discipline protects. The United Nations Universal Declaration of Human Rights recognizes privacy as a fundamental right, and many jurisdictions codify it in law. One of the primary drivers of cybersecurity *is* to support privacy by limiting access to private information. Yet a myth persists that increasing privacy protections necessarily reduces a system's security, and vice versa. As a general claim, that is false. That myth arises because the cheapest solution to some security problems involves reducing privacy. For example, an organization may store copies of all incoming email to detect phishing, even though other methods, such as automatically rewriting URLs to neutralize harmful links, require neither surveillance nor recording of email content. The tension is real in some other cases, however. Monitoring for insider threats, forensic retention of logs, and lawful-intercept obligations each trade some privacy for some security, and no design trick makes the trade disappear. What a definition can do is make the trade visible. If privacy is among the interests being protected, then a control that erodes it carries a security cost, to be weighed under the governing policy rather than treated as free. People who work in cybersecurity should protect privacy, not erode it. Faced with regulatory requirements such as those of the GDPR, practitioners have a professional duty to support them rather than circumvent them.

Safety is an adjacent discipline that overlaps with cybersecurity rather than being part of it. Software now mediates safety-critical functions in medical devices, vehicles, industrial control systems, and critical infrastructure. Safety engineering concerns harm to people and property from every cause. Cybersecurity concerns harm mediated through computing-dependent systems. Convention separates safety from security by intent, with accidental harm on one side and deliberate harm on the other. The definition proposed here does not use intent as its boundary, for two reasons. First, practice does not draw the line there, as the outage case below shows. Second, intent is often indiscernible. A crashed controller looks the same in the first hours whether a bug or an attacker caused it, and attribution, when it comes at all, comes after the response has been initiated. A definition that depends on intent cannot say what an incident is until the investigation is over, if then. The definition uses the medium instead and accepts that overlap may occur. Where the two overlap, such as in cyber-physical systems, a cybersecurity failure is a safety failure, and the reverse may also hold. A flawed patch to an industrial control system or medical device may introduce failures that are not security failures, *per se*, but are as dangerous as the vulnerability it was

meant to address. A definition of cybersecurity should therefore count harm to safety among the harms it manages, without claiming the whole of safety engineering as its own.

Ethics differs from both because it governs the practice, not the system. As a *property* of a system, security has no ethical content, because a system either behaves according to its policy or it does not. As a *discipline*, cybersecurity is a professional practice, and definitions of professional practices often say whom the practice serves. ABET's long-standing definition of engineering, inherited from its predecessor, the Engineers' Council for Professional Development, ends with the purpose "for the benefit of mankind."[c] The IAEA defines safety as protection "of people and the environment." Those are purpose clauses that name a beneficiary, and this article's thesis supplies one. Security exists to protect the primary missions of the people and organizations that depend on computing. That purpose clause alone is not enough, however, because a practice can serve its own beneficiaries at everyone else's expense. A ransomware gang may manage risk to its own systems relative to a context and a policy, possibly with considerable technical skill. Its systems may well be secure in the property sense. That practice is not cybersecurity in the discipline sense this article intends, and only a normative element in the definition says why. Professional codes of ethics, such as those of the ACM, the IEEE, and the Information Systems Security Association (ISSA), already spell out the duties: care for society and human well-being, responsibility for ethical practice, and respect for privacy. A cybersecurity definition need not enumerate them, and it should not invite argument over whose ethics apply. It need only say that the profession's own codes bound the practice. Without that clause, it describes a skill set rather than a profession.

Vulnerability disclosure illustrates the point. No U.S. law requires a researcher who finds a critical flaw in a widely used system to tell the vendor before publishing it. The norms that fill the gap, from coordinated disclosure to fixed deadlines, are contested among practitioners, and automated discovery strains all of them. When tools find flaws faster than vendors can fix them, no single schedule fits every case. What remains is judgment about harm to the people who use the system, which is what professional ethics governs. That we *can* do something does not mean we *should*, and a definition of the discipline should reflect that.

A Proposed Definition, and Some Hard Cases

With those requirements in mind, here is a formulation, offered as a starting point for discussion:

> *Cybersecurity is the ongoing management of risk to computing-dependent systems, and to the people, organizations, and missions that rely on them, relative to context and policy and within the bounds of professional ethics, spanning technical, human, organizational, and societal dimensions.*

Each clause has a purpose. *Ongoing* rejects the end-state fallacy. Security is not achieved and then held. *Management of risk*, rather than threat elimination, aligns with Courtney's and Schneier's framing and makes the definition compatible with quantitative methods. The risk in question is that a system, or the people and processes around it, will behave contrary to the governing policy, from whatever cause. This is deliberately broader than "unauthorized." A disk that fails without a backup violates an availability requirement as surely as an intruder who deletes the data, and the 1991 formulation quoted earlier,

[c] Approved by the ECPD Board of Directors in 1979, extending a 1961 formulation; see National Research Council, *Engineering Infrastructure Diagramming and Modeling*, Appendix A (Washington, DC: National Academies Press, 1986), https://nap.nationalacademies.org/read/587/chapter/8.

"behave as expected," already took that view. How an organization classifies such an event is operationalization. *Computing-dependent systems* is deliberately broader than "information systems" or "computers and networks." It includes data, software, hardware, and networks, as well as the cyber-physical devices, vehicles, medical instruments, and industrial controllers whose behavior depends on computation. *The people, organizations, and missions that rely on them* places the primary goals discussed under "Security as Instrument, Not End" at the center, and it brings social engineering, fraud, privacy, and safety into scope, because in each case the harm falls on people through the systems they depend on. *Relative to context and policy* is Courtney's first law, and it is also where legitimate variation lives, because policy is set locally. *Within the bounds of professional ethics* qualifies the discipline, not the property, and names the obligation to those the discipline serves, for the reasons given above. The definition covers both senses of the word. The *property* is the state in which such risk is managed to the level the policy accepts, and the *discipline* is the practice of reaching and maintaining that state.

A definition earns confidence by handling cases that strain it. Five follow, chosen because each exposes a boundary where existing definitions disagree.

Deepfake fraud. In February 2024, attackers used AI-generated video and audio of a company's executives to conduct a fraudulent video conference, and an employee transferred $25 million to them. No system is reported to have been penetrated, and the videoconferencing software worked as designed. Under Merriam-Webster's definition, and on a plausible reading of the SEC's "unauthorized occurrence," this is a fraud rather than a cybersecurity incident. Under the proposed definition it is squarely in scope because the attackers used a computing-dependent system against the people and the mission that relied on it. The definition does not decide whether the victim's cyber insurance policy or its crime policy pays. That is operationalization, and it belongs to the contracts.

A non-malicious outage. On July 19, 2024, a faulty configuration update to CrowdStrike's Falcon security software crashed roughly 8.5 million Windows systems, grounding flights and disrupting hospitals and banks.[d] Nobody attacked anything. Definitions built on an adversary or an unauthorized act, such as Merriam-Webster's and the SEC's, exclude the event. RFC 4949's, which counts "accidental change, destruction, or loss," and the FSB's Cyber Lexicon, which deliberately includes non-malicious events, admit it. So does the proposed definition, because the risk was to computing-dependent systems and to the missions that relied on them, and the source, a security product, is immaterial.

Monitoring versus privacy. Suppose an organization proposes to retain every inbound email for a year to improve phishing detection. Under the proposed definition, privacy is among the interests protected, so the retention is a security cost as well as a security benefit, and the policy must weigh the two. The definition does not say which way the balance falls. A hospital and an intelligence agency will decide differently, and both may be right. It does say that the question must be asked, and that the ethical clause limits the acceptable answers.

A cyber-physical patch. Suppose a vendor issues a security patch for an infusion pump. The patch closes a remote-access vulnerability but introduces a timing fault that, under rare conditions, halts the pump. The vulnerability was a cybersecurity problem according to any definition. Is the patch's side effect also one? Under definitions confined to computers and information, no, because it is a safety defect. Under the proposed definition, yes, because the patch was a risk-management action on a

[d] Microsoft's estimate; see D. Weston, "Helping our customers through the CrowdStrike outage," Microsoft, July 20, 2024, https://blogs.microsoft.com/blog/2024/07/20/helping-our-customers-through-the-crowdstrike-outage/.

computing-dependent system, and it created new risk to the people relying on it. Safety engineering owns the hazard analysis, but cybersecurity owns the decision to patch and its consequences.

Artificial intelligence. AI is the sharpest test of technology neutrality, because it changes attack and defense at the same time. Large language models generate personalized phishing at scale, automated tools discover and exploit vulnerabilities faster than human defenders can patch, and AI systems defend networks with increasing autonomy, which makes the defender itself an attack surface. Prompt injection against an AI agent is recognizably an injection attack, a class the field named decades ago when untrusted input was first mistaken for instructions, which is itself evidence that the vocabulary transfers. The proposed definition names no technology, so none of this stretches it. An AI agent that acts on behalf of a user is a component of a computing-dependent system whose behavior must be managed relative to policy, whether it defends, attacks, or misbehaves on its own. The risk it introduces is judged by the same standard as any other. If a definition survives the arrival of a technology that automates both attack and defense, it should survive the next one.

The definition also has limits. A disinformation campaign conducted through accounts used in accordance with a platform's rules harms people through a computing-dependent system, yet the system behaved as its policy specifies. By the rule stated above, that is not a cybersecurity problem, however serious a problem it is. The same campaign conducted through compromised or fraudulent accounts is one. The clause on professional ethics excludes ransomware gangs' operational security from the discipline, no matter how secure their systems are as a matter of property. The definition does not resolve every boundary dispute, and it is not intended to. Instead, it provides a consistent way to consider one.

The Case for Convergence

The field's definitional failure is not merely pedantic. It has practical consequences for metrics development, comparison of defensive approaches, education, regulatory coherence, public trust, workforce estimates, and risk underwriting.

A Common Core, Not Uniformity

The definition proposed above is one candidate for a shared core, but there could be others. The argument does not depend on its particular wording but on the shape of the objective: a common conceptual core, precise enough to anchor metrics and cross-jurisdictional comparison, with context-specific operationalization left to sectors and jurisdictions. That is how the comparator fields work. ICAO's "acceptable level" of risk is the shared referent, while each state sets thresholds for its own airspace. Complete uniformity is neither achievable nor desirable, because a hospital, a bank, and a power utility face different risks and different regulators, and their operational definitions of "incident," "significant," and "secure enough" should differ. Variation of that kind is productive. Variation in what the word refers to is not, because it makes the productive kind invisible.

A converged definition of cybersecurity likewise need not specify which threats matter, which technologies to deploy, or which controls to require. It need not resolve every boundary dispute in advance. The hard cases above show how to evaluate one. It need not be static. All three comparator definitions have been revised, and given the tempo of adversarial adaptation, a definition of cybersecurity should carry explicit revision cycles from the start. What it must do is provide a shared referent, precise enough to anchor metrics, broad enough to accommodate context, and stable enough to support cross-

jurisdictional comparison: a common answer to the question "What are we talking about when we say cybersecurity?"

An Institutional Pathway

Calling for definitional convergence without identifying a realistic institutional pathway is merely hand-wringing. The analogies above are instructive. Each succeeded because a single body with treaty-level authority convened the effort and had the standing to drive adoption. Cybersecurity has no comparable institution.

Several candidates exist, each with limitations. ISO/IEC JTC 1/SC 27 has the broadest international membership. It has already produced the ISO/IEC 27000 family, but it declined to define "cybersecurity" in ISO/IEC 27032:2023, opting instead for cross-references. NIST has enormous influence, particularly in the United States, but its definitions are not binding on other nations, and its six internal definitions of cybersecurity illustrate the problem rather than solving it. The EU, through ENISA and the NIS2 Directive [4], has demonstrated willingness to legislate definitions, but EU regulation applies only within its jurisdiction.

A more promising model may be a coalition in which two or three of these bodies produce a definition endorsed simultaneously by ISO, NIST, and ENISA (or their successors), with buy-in from sector-specific regulators such as the Financial Stability Board for finance and ICAO for aviation. The key is that someone must convene the effort, fund it for the long haul, and commit to iterating through what will inevitably be a contentious process. Half-measures—yet another working group that produces yet another incompatible definition—will make the problem worse, not better. Any such effort must also include meaningful participation from developing nations and civil society. Without their voices, the result will lack the legitimacy that broad adoption requires.

The Cost of Continued Inaction

The cost of definitional ambiguity resists precise estimate, and this article offers none. The scale of the activity affected is not in doubt. Gartner forecast worldwide end-user spending on information security to exceed $200 billion in 2025.[e] If definitional ambiguity introduces even modest inefficiencies, the aggregate waste would be large in absolute terms, though no one can say how large. The GAO's finding that insurers cannot consistently define the risk they underwrite [7] means that insurers price premiums on a definitionally unstable basis, and misdefined risk means mispriced coverage. The workforce numbers deserve the same caution. A broad definition produces a larger gap than a narrow one, so some fraction of the widely cited shortfall is an artifact of measurement rather than of demand. How large a fraction is unknown, and no one should claim otherwise. These numbers cannot be made precise because the definitions they depend on are unsettled. That is the point.

Call to Action

Three related misconceptions persist: that we know what cybersecurity is, that it is an achievable end state, and that existing technology is sufficient. Recognizing these misconceptions is the first step toward addressing them. The second step is a serious, multidisciplinary definitional effort that draws on computing, risk analysis, organizational behavior, ethics, law, and policy to establish a shared core on which sectors and jurisdictions can build. It will not be easy, but neither was harmonizing safety

[e] Gartner, "Gartner Forecasts Worldwide End-User Spending on Information Security to Total $213 Billion in 2025," press release, July 29, 2025.

definitions for an industry in which hundreds of airlines operate in 193 sovereign nations. The aviation community managed it because the cost of *not* having a shared definition was too high. We are approaching that threshold in cybersecurity.

Aristotle observed that it is the mark of an educated mind to look for precision in each class of things only so far as the nature of the subject permits [14]. Cybersecurity does not need a perfect definition, only one that is precise enough to ground measurement and broad enough to encompass the full scope of what cybersecurity must protect. Until we have one, we will keep building artifacts that embody uncertain features interacting in surprising ways, including potentially catastrophic failures.

Acknowledgment

This article adapts and significantly extends material from Chapter 1 of *Cybersecurity Myths and Misconceptions: Avoiding the Hazards and Pitfalls that Derail Us* by G. Spafford, L. Metcalf, and J. Dykstra (Addison-Wesley, 2023) [15]. The author was the primary contributor to that chapter. Used with permission of the publisher and co-authors. The author thanks Leigh Metcalf and Josiah Dykstra for their contributions to the original material and for their encouragement on this adaptation, and the anonymous reviewers for comments that sharpened the argument.

Portions of this text were edited and structured using Anthropic Claude (Opus 4.6 and Fable 5.1) based on the author's drafts; the ideas, arguments, and final editorial decisions were the author's, as are the em dashes and semicolons.

References

[1] National Institute of Standards and Technology, "CSRC glossary: Cybersecurity," https://csrc.nist.gov/glossary/term/cybersecurity, 2024, accessed: 2026-04-06.

[2] S. Garfinkel and G. Spafford, *Practical UNIX Security*. Sebastopol, CA: O'Reilly & Associates, 1991.

[3] "Regulation (EU) 2019/881 of the European Parliament and of the Council on ENISA and on information and communications technology cybersecurity certification," Official Journal of the European Union, 2019, article 2, point (1). [Online]. Available: https://eur-lex.europa.eu/eli/reg/2019/881/oj/eng

[4] "Directive (EU) 2022/2555 on measures for a high common level of cybersecurity across the Union," Official Journal of the European Union, 2022. [Online]. Available: https://eur-lex.europa.eu/eli/dir/2022/2555

[5] D. Craigen, N. Diakun-Thibault, and R. Purse, "Defining cybersecurity," *Technology Innovation Management Review*, vol. 4, no. 10, pp. 13–21, 2014. [Online]. Available: https://timreview.ca/article/835

[6] Organisation for Economic Co-operation and Development, "New perspectives on measuring cybersecurity," OECD, Digital Economy Papers 366, Jun. 2024. [Online]. Available: https://www.oecd.org/en/publications/new-perspectives-on-measuring-cybersecurity_b1e31997-en.html

[7] U.S. Government Accountability Office, "Cyber insurance: Insurers and policyholders face challenges in an evolving market," GAO, Tech. Rep. GAO-21-477, May 2021. [Online]. Available: https://www.gao.gov/products/gao-21-477

[8] *Annex 19 to the Convention on International Civil Aviation: Safety Management*, 1st ed. Montréal, QC, Canada: Int. Civil Aviation Org., Jul. 2013. [Online]. Available: https://www2023.icao.int/NACC/Documents/Meetings/2014/SSPSMS/Annex19.pdf

[9] International Atomic Energy Agency, "IAEA nuclear safety and security glossary," 2022, interim 2022 edition. [Online]. Available: https://www-pub.iaea.org/MTCD/Publications/PDF/IAEA-NSS-GLOweb.pdf

[10] International Organization for Standardization, "ISO 9000:2015 — quality management systems — fundamentals and vocabulary," 2015. [Online]. Available: https://www.iso.org/standard/45481.html

[11] S. J. Choi, M. E. Johnson, and C. U. Lehmann, "Data breach remediation efforts and their implications for hospital quality," *Health Services Research*, vol. 54, no. 5, pp. 971–980, 2019.

[12] B. Schneier, "Crypto-gram newsletter," https://www.schneier.com/crypto-gram/archives/2000/0515.html, 2000, may 15, 2000.

[13] W. D. Young, W. E. Boebert, and R. Y. Kain, "Proving a computer system secure," *Scientific Honeyweller*, vol. 6, no. 2, pp. 18–27, Jul. 1985.

[14] Aristotle, *Nicomachean Ethics*. Princeton, NJ: Princeton University Press, 1984, trans. W. D. Ross, rev. J. O. Urmson, in J. Barnes, Ed., *The Complete Works of Aristotle*, vol. 2, Book I, Ch. 3, 1094b.

[15] G. Spafford, L. Metcalf, and J. Dykstra, *Cybersecurity Myths and Misconceptions: Avoiding the Hazards and Pitfalls that Derail Us*. Boston, MA: Addison-Wesley Professional, 2023.

Eugene H. Spafford is a Distinguished Professor of Computer Science and founder and Executive Director Emeritus of CERIAS at Purdue University, West Lafayette, IN 47907 USA. His research interests include cybersecurity, cyber policy, and computing ethics. Spafford received his Ph.D. in information and computer science from the Georgia Institute of Technology. He is an ACM Fellow and Life Member, an IEEE Life Fellow, a Fellow of the AAAS, a Distinguished Fellow of the ISSA, and a member of the Cyber Security Hall of Fame. Contact him at spaf@purdue.edu.