For Good Measure

The Imperative of Reclaiming Metrics Terminology

DAN GEER AND JASON CRABTREE

Dan Geer is a Senior Fellow at In-Q-Tel and a security researcher with a quantitative bent. He has a long history with the USENIX Association, including officer positions, program committees, etc. dan@geer.org

Jason Crabtree is the CEO and co-founder of QOMPLX, with a focus on cybersecurity, operational risk management, and decision support technology. Prior to launching QOMPLX, he served as Special Advisor to the Commanding General of Army Cyber Command, as an infantry leader in Afghanistan, and holds degrees from West Point and Oxford University, where he studied as a Rhodes Scholar. jason@qomplx.com

I'm swimming
in darkness
keeping eyeballs clear
—Murio Suzuki
(Trans. Ban'ya Natsuishi)

The explosion of interest in measuring and reporting on security has been most welcome, yet that surge has also brought with it powerful side effects, often stemming from a lack of consistent ontology to aid in common understanding, reasoning, and communication. Many current efforts suffer from a misunderstanding of the distinct differences between data, information, knowledge, and wisdom. We are too often speaking past one another—even more so as information technology, business, legal, and other professions collide.

Our primary purpose in the field of risk management must be to improve future outcomes for our stakeholders. The requisite discipline required—to perpetually focus on this goal and to avoid the siren call to seek ever higher fidelity of retrospective justifications for after-the-event opinions with which to blame or litigate one another—is substantial.

With this in mind, it is worth revisiting several central tenets which support this focus on *ex ante* decision-making against which we should hold individuals and organizations accountable versus *ex post* claims of negligence or the too-often hypothetical "we could have done *X* to prevent this."

> "When you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge is of a meager and unsatisfactory kind: it may be the beginning of knowledge, but you have scarcely, in your thoughts, advanced to the stage of science."—*William Thomson, Lord Kelvin*

First, we seek to leverage quantitative and qualitative measures of risk in order to support our own internal reasoning but ultimately to support our collective reasoning and interactions. Encouraging accountability and economically rational actions in a complex multi-agent decision-making environment demands semantically consistent approaches. This is at the core of linking tactical operational security decision-making with enterprise-level risk management with supply chain and counterparty risk management with policymakers', regulators', and economists' actions. Said more poetically, without a consistent ontology, "Meaning lies as much in the mind of the reader as in the Haiku" [1].

Central Thesis

The central thesis of this essay is so aligned: a sufficient amount of activity around the concept of cyber risk without the requisite degree of specificity or consistency is masking a lack of sufficient, fundamental progress in the true science Kelvin implores practitioners

to seek. The result is confusion of activity with achievement. When information becomes cheap, attention becomes expensive, and our rapid instrumentation of enterprise networks and the broader Internet has yielded a wave of information with equal parts utility and distraction.

A potent illustration of the growing phenomenon is the overstatement of individual metrics or groups of metrics to comment on the security of individual organizations or groups of organizations, or to characterize broader systemic risk, e.g., for financial services or utilities, based on myopically focused collections of numbers and tenuous correlations to poorly sampled breach or loss events. Examples include:

1. Misuse of CVSS scores for vulnerability patching and prioritization efforts
2. Misrepresentation of external scan data as a proxy for holistic security posture
3. Lack of accurate characterization about TCP/UDP DDoS vuln and bot activity (see, e.g., when a Fortnite update was anomalous enough that assertions of DDoS attacks were thrown around as a result of insufficient correlation between system perturbations and environmental changes which are larger than historical baseline model anomalies) [2]
4. Insufficient research into BGP protocol issues
5. Use of behavioral analytics to claim comprehensive insider threat modeling despite widespread forging and manipulation of Kerberos SSO or even vendors claiming that they can "secure" fundamentally insecure protocols like NTLM with multi-factor authentication and heuristics

In some ways our issues revolve around our lacking the ability to understand the value of information remaining confidential, retaining its integrity, or being available. Add in the value of that same information being presented at the right time, in the right place, with sufficient context and we have captured our collective challenge as practitioners of operational risk management—something well beyond cybersecurity alone. That portion of the problem remains out of scope here.

Examples Appear

The appearance of larger limits and now larger resultant losses in cyber insurance is instructive. Global insured losses from NotPetya and other ransomware attacks on a claims-made basis have reached more than $3B in aggregate—with around 90% driven by silent cyber impacts and the remainder from affirmative losses to specific cyber insurance contracts [3]. Economic losses exceeded $10B in total [4].

Digging into some of the litigation underscores the importance of definitions of terms/entities and the ability to manage large amounts of data associated with determining whether specific facts can be supported via available information and whether or not specific aspects of the contracts relating the different counterparties are impacted by those facts.

A major company, Mondelēz, claimed $100 million on its insurance policy because it believed the permanent damage to 1,700 servers and 24,000 laptops, theft of thousands of user credentials, business interruption, and lost revenue from unfulfilled customer orders were compensable under the provision of an insurance policy that covered "physical loss or damage to electronic data, programs, or software" caused by "the malicious introduction of a machine code or instruction" or from the failure of Mondelēz's electronic data processing equipment or media. Zurich's counter that no payment was due as a result of an exclusion for "hostile or warlike action in time of peace or war" has led to litigation [5]. Tracking the percentage of cyber-insurance events and policies that lead to litigation may prove to be a proxy for tracking the degree to which there is misalignment between technical, business, and legal considerations.

The confusion about terminology and even how the courts may interpret such language is impactful. Regulated financial institutions and other industries who have specific capital requirements use insurance products to transfer risk off of their balance sheets, but this type of litigation undermines confidence for risk managers and regulators that such capital will be paid out in a timely fashion; this, in turn, exacerbates basis risk and potentially makes certain insurance policies incompatible with broader regulatory capital requirement wording requirements [6].

Our dependence on all things cyber as a society is now inestimably irreversible and irreversibly inestimable. Since dependence (and interdependence) continue to grow, we cannot understand the ordinate values, but we can understand the trend and the degree to which select risks are convex or concave.

In Comparison Is Insight

Even if an organization is able to internally capture and correlate its operational disruptions or losses to various metrics, without a consistent ontological perspective to share among its peers, it is not possible to robustly understand or track changes in systemic risk. Again, if all organizations have somewhat similar ideas of a set of metrics and generally believe themselves to be experiencing the same convex (e.g., DDoS attacks) or concave trend (e.g., falling price of stolen financial system identities/records in absolute terms or as normalized against health-care records), then some conclusions may be drawn. However, if there are differing perspectives (especially within peer groups with a high degree of similarity), then new challenges arise.

Systemic risk analysis, which by definition is incorporating data from multiple entities, also requires better insight into ordinality than self-referential comparisons within a single organization.

While staff do change, in general most larger institutions have an established culture associated with the process for data collection, analysis, and reporting that enables some consistency, however imperfect. The lack of reference scenarios for calibration purposes, ontologies for a common entity, and even field mapping is problematic. That said, techniques like Business Process Management and Notation (BPMN) and universal metric types, e.g., mean time between failures (MTBF) and mean time to repair (MTTR), can help. If the process-centric BPMN definitions are combined (and harmonized) with concepts contained in other developing standards such as MITRE ATT&CK (for threat tactics, techniques and procedures; https://attack.mitre.org), STIX2.0 (for threat intelligence/actor data; https://oasis-open.github.io/cti-documentation/), and OGIT (for asset data; https://github.com/arago/OGIT/), then more meaningful excavation of relationships between assets, processes, impacts, and actions from internal staff or external threat actors is possible.

Design scenarios provide useful validation mechanisms for a broader ontological design process, but also enable individual teams and organizations to translate their internal efforts into a more universally communicable framework. One exemplary tool which should be considered is the Cambridge Center for Risk Studies' taxonomy of business risks, which is being improved to capture key aspects of cyber events and technology risks more broadly [7]. If coupled with better disclosure from all parties, we can do a better job of understanding the relationships between business impacting events, financial losses, and the actual specifics of various accidental failures or targeted incursions.

Comparison Requires Communication

We often note that people reason by analogy and the common lazy cyber analogies of soccer, war, etc., end up being misused as a direct result of the same lack of specificity in the underlying ontologies and scenarios for individual problem representation and transformation. Metric communication about the appropriate trends, ordinal elements, and links of those metrics to specific assets and processes of material interest to leadership and customers (or consuming such data from suppliers when considering third- and fourth-party risk) depend on the ability to tell a story. These scenario-based narratives enable us to connect general structure with specific instances where individual people and organizations have direct familiarity. "This idea that there is generality in the specific is of far-reaching importance" [1].

Take, for example, the Basel Committee on Banking Supervision's definition of a risk concentration as "an exposure with the potential to produce losses large enough to threaten a financial institution's health or ability to maintain its core operations" [8]. The lack of a sufficiently generalized reference model for data and scenario capture precludes efficient or consistent evaluation of any given portfolio of metrics. For example, if there is no shared ontology for users, hosts, privileges, network topologies, and business processes and their relationships, it becomes virtually impossible to make useful comparisons across more than one entity even if they were simplified and we pretended that technology, defender behavior, and attacker capabilities were static.

The gaps in current approaches become even more apparent when attempting to capture elements of the learning inherent in battles between sentient actors with their own economic constraints. Simply put, reasonably modeling non-random (non-ergodic) agent and system behavior requires correlation of business-impacting events and losses in a rigorous fashion. It requires a keen understanding of internal, external, Internet infrastructure, threat actor/geopolitical, environmental conditions, and more—it is not a simple retrospective modeling task.

We know this from other forms of risk modeling, particularly around crisis modeling, where, by definition, events are not particularly similar to the past and initial shocks can lead to cycles of behavior that reverberate across local and global incentives and decision constraints practically imposed on other actors.

"Discovering vulnerability to crisis requires a specification of system dynamics and behavior. Even if we are willing to make the leap of asserting that any one financial institution is not large enough for a stress to affect other parts of the financial system, if banks share similar exposures and thus are affected similarly by the stress, the aggregate effect will not be likely to reside in a ceteris paribus world. Furthermore, in the highly interrelated financial system, the aggregate effect will feed into yet other institutions and create adverse feedback and contagion" [9].

Key definitions of terms and agreement on real ontological frameworks cannot be left to the flamboyant misappropriation and misuse of terms like "resilience" in the press and in marketing material. These terms have value and they are central to meaningful communication about our individual, organizational, sector, and broader societal exposure to dependence on technologies, common infrastructure, and one another. Our growing exposure to transitive risks associated with interdependence demands robust efforts to set the stage for collaboration around metrics—which starts with doing the difficult work of ontology specification.

No one said this would be easy.

References

[1] D. R. Hofstadter, *Gödel, Escher, Bach: An Eternal Golden Braid* (Basic Books, 1979).

[2] Stilgherrian, "Suspected Commonwealth Games DDoS Was Only a Fortnite Update," *ZDNet*, September 11, 2019: https://www.zdnet.com/article/suspected-commonwealth-games-ddos-was-only-a-fortnite-update/.

[3] L. Gallin, "NotPetya Insured Loss Could Creep 30%+ as Tail Develops: Johansmeyer, PCS," *Reinsurance News*, August 14, 2019: https://www.reinsurancene.ws/insured-notpetya-loss-could-creep-30-as-tail-develops-johansmeyer-pcs/; and S. Evans, "Mondelēz's NotPetya Cyber Attack Claim Disputed by Zurich," *Reinsurance News*, December 17, 2018: https://www.reinsurancene.ws/mondelezs-notpetya-cyber-attack-claim-disputed-by-zurich-report/.

[4] A. Greenberg, "The Untold Story of NotPetya, the Most Devastating Cyberattack in History," *Wired*, August 22, 2018: https://www.wired.com/story/notpetya-cyberattack-ukraine-russia-code-crashed-the-world/.

[5] L. Bershidsky, "Zurich Policyholder Dispute Highlights Danger of Calling Out Cyber Attackers," Insurance Journal, January 11, 2019: https://www.insurancejournal.com/news/international/2019/01/11/514553.htm; and "Mondelēz Sues Zurich in Test for Cyber Hack Insurance," *Financial Times*, January 11, 2019: https://www.ft.com/content/8db7251c-1411-11e9-a581-4ff78404524e.

[6] A. Satariano and N. Perlroth, "Big Companies Thought Insurance Covered a Cyberattack. They May Be Wrong," *The New York Times,* April 15, 2019: https://www.nytimes.com/2019/04/15/technology/cyberinsurance-notpetya-attack.html.

[7] A. Coburn, "The Future of Cyber Risk," Cambridge Centre for Risk Studies, July 2019: https://www.jbs.cam.ac.uk/fileadmin/user_upload/research/centres/risk/downloads/192407_cyberconference_presentation_coburn.pdf.

[8] The Joint Forum, "Risk Concentration Principles," 1999: https://www.bis.org/publ/bcbs63.pdf.

[9] R. Bookstaber, M. Padrik, B. Tivnan, "An Agent-Based Model for Financial Vulnerability," Office of Financial Research, US Treasury, September, 2014: https://www.financialresearch.gov/working-papers/files/OFRwp2014-05_BookstaberPaddrikTivnan_Agent-basedModelforFinancialVulnerability_revised.pdf.