Evidence-Based Prioritization of Cybersecurity Threats

The landscape of cybersecurity threats is evolving at a tremendous pace. New adversaries with increasingly sophisticated tactics and tradecraft seem to emerge on a continuous basis. Meanwhile, the information and communications technology (ICT) infrastructures being targeted by these adversaries are becoming more diverse and complex, not least because of the rapid deployment of new technologies such as the cloud and the Internet of Things (IoT). To stay abreast of both imminent and emerging threats, many enterprises have invested in so-called threat intelligence capabilities. By structurally collecting threat-related data (e.g., concerning the objectives and tradecraft of adversaries targeting their industry), such enterprises pursue situational awareness and strive to anticipate upcoming threats rather than waiting for an actual incident to occur. For instance, new insights or notable trends might trigger a change in technical infrastructure or the targeted education of security staff to prepare for a particular (new or evolving) threat. The intelligence that enterprises collect to this end comes in a variety of forms, but it is increasingly structured (machine readable) in nature and is often fed to a dedicated threat intelligence platform (TIP) for storage and processing.

Enterprises cannot take proactive precautions for every cybersecurity threat that comes to their attention, so it is imperative that they set appropriate priorities, typically by maintaining a so-called threat landscape that ranks threats in order of importance. Currently, such priorities are often

Richard Kerkdijk
Is a senior security consultant at the Netherlands Organisation for Applied Scientific Research (TNO). His role involves strategic advisory work, technical and nontechnical security evaluations, and coordination of cybersecurity research and innovation projects. He deals mostly with telecommunications providers (across Europe) and financial institutions (in The Netherlands), but he also has done work for the Dutch National Cyber Security Center, the Dutch Cyber Security Council and the Dutch Ministry of Defense. He is vice chair of the ETIS Information Security Working Group, an industry body that facilitates collaboration among chief information security officers (CISOs) of European telecommunications providers.

Sebastiaan Tesink, CISA, CISM, CISSP
Is a security researcher in the Cyber Security and Robustness Department at the Netherlands Organisation for Applied Scientific Research (TNO). His work focuses on the automation of security operations centers (SOCs) and computer security incident response teams (CSIRTs) and vulnerability research.

Frank Fransen
Is a senior scientist in the Cyber Security and Robustness Department at the Netherlands Organisation for Applied Scientific Research (TNO). His work involves the study of emerging security technologies, security of mobile communication systems (3G, 4G and 5G), information security and risk management, security operations, cyberthreat intelligence, and cybersecurity of smart energy grids.

Federico Falconieri
Is a junior cybersecurity specialist in the Cyber Security and Robustness Department at the Netherlands Organisation for Applied Scientific Research (TNO). He works on security automation projects such as development, security and operations (DevSecOps); unsupervised and semisupervised network attack detection; threat intelligence enrichment and distribution; and automated incident response.

driven by an enterprise's security practitioners and their expert appraisal of the relevance and severity of specific threats. This raises issues of credibility and accuracy, especially if the recommended precautions will have a strong impact on the enterprise's business or if they involve major investments. To strengthen the foundation for strategic security decision making, it is preferable to prioritize threats on the basis of actual observations (evidence) rather than human opinions. This can be achieved by leveraging the vast amount of threat-related data that enterprises maintain in TIPs and operational tools, such as security monitoring and incident workflow solutions. The model presented converts such data into quantitative (metrics-driven) scores that reflect the priority of threats for individual enterprises or for the broader industry in which they reside.

Building a Model

To prioritize threats, the first step is to ask the question: What is a threat? Although this may seem obvious, the topics that threat reports include under this heading can vary from threat actors (e.g., insider threats) and campaigns (e.g., cryptojacking, cyberespionage) to attack techniques (e.g., malware, phishing) and even general technology trends (e.g., cloud, IoT, privacy and data protection).[1, 2] Although a mixture of concepts may be appropriate for particular threat landscapes, prioritizing threats for an enterprise or industry requires more specificity and consistency.

> "TO STRENGTHEN THE FOUNDATION FOR STRATEGIC SECURITY DECISION MAKING, IT IS PREFERABLE TO PRIORITIZE THREATS ON THE BASIS OF ACTUAL OBSERVATIONS (EVIDENCE) RATHER THAN HUMAN OPINIONS."

In general, a threat is anything (e.g., object, substance, human) that is capable of acting in a manner that can result in harm.[3] Thus, the term "threat" essentially refers to the threat actors (or threat agents) who might target an enterprise (it can also refer to natural causes, but those are not the focus here). Appropriately ranking the extent to which such actors can cause harm requires a consideration of their intention. In cyberthreat intelligence terms, this is typically referred to as the campaign of the threat actor, which can be defined as a set of adversarial behaviors that constitute a set of malicious activities or attacks occurring over a period of time against specific targets.[4] It is therefore assumed that an enterprise needs to assign priorities to campaigns conducted by particular actors.

The essential goal of threat prioritization is to determine the likelihood that a specific campaign (carried out by a particular threat actor) will manifest itself within the enterprise under consideration. This likelihood should be deduced from actual threat events observed by the enterprise (e.g., events in its security monitoring systems) or retrieved from external sources (e.g., public threat reports, commercial cyberthreat intelligence [CTI] feeds, closed CTI communities). Typically, risk assessment methodologies do not provide guidance or models to weigh these threat events and make a quantitative appraisal of the risk they pose. However, Factor Analysis of Information Risk (FAIR) provides a model for understanding, analyzing and quantifying cyberrisk that includes a taxonomy of factors contributing to such risk (**figure 1**).[5, 6] The loss event frequency (LEF) leg of the model provides a particularly good starting point for evidence-based threat prioritization.

LEF is defined as the probable frequency within a given time frame with which a threat agent will inflict harm on an asset. The LEF parameter depends on the probability that a threat actor will act against an asset—the threat event frequency (TEF)—and the probability that the threat actor's actions will be successful (vulnerability). For the purpose of prioritizing threats, LEF is the probable frequency with which the threat actor will

Figure 1—FAIR Risk Taxonomy Focusing on Loss Event Frequency

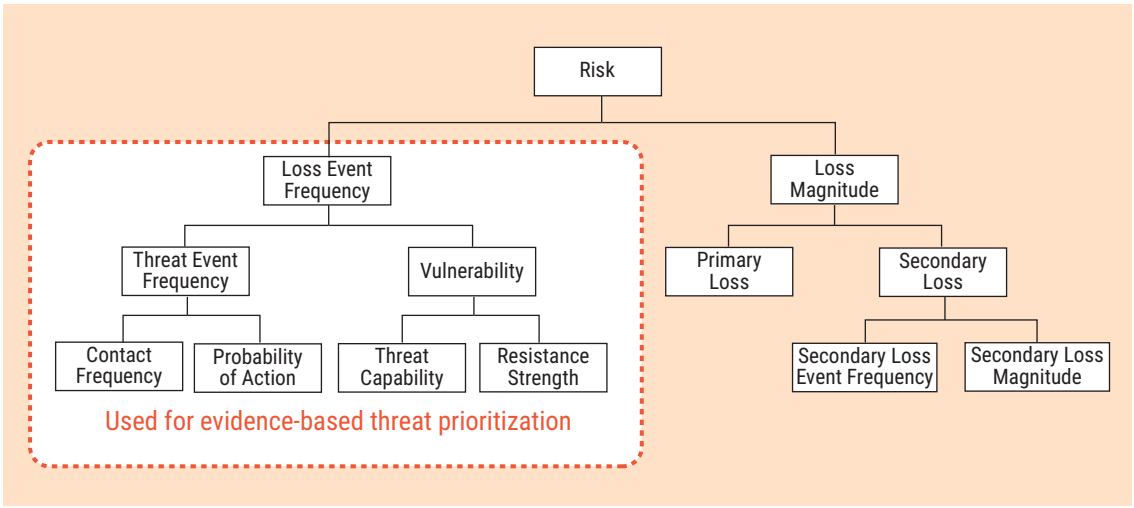

successfully execute the campaign against the enterprise under consideration.

The TEF parameter is made up of two factors:

1. **Contact frequency**—The probable frequency with which a threat actor will come into contact with an asset
2. **Probability of action**—The probability that a threat agent will act against an asset once contact has occurred

For the purpose of prioritizing threats, the contact frequency is essentially the likelihood that an organization will show up on the threat actor's radar, and the probability of action is the likelihood that the threat actor will actually target the organization as part of an ongoing campaign.

The vulnerability parameter also comprises two distinct factors:

1. **Threat capability**—The probable level of force a threat agent is capable of applying against an asset.
2. **Resistance strength**—The strength of a control compared with a baseline measure of force.

For purposes of evidence-based threat prioritization, these two factors can be combined by mapping the tactics, techniques and procedures (TTPs) the threat actor has employed in this or similar campaigns to the presence and strength of the security controls implemented by the organization. The idea is to assess which techniques and procedures the actor can and cannot successfully execute. Other factors may also be considered, such as the threat actor's resources and adaptability (threat capability) and the patch levels of the organization's technical assets (resistance strength).

Specifying Threat-Oriented Metrics

To create an auditable evidence-based threat prioritization, the FAIR factors must be divided into observable and measurable elements—evidence-based threat categories—that can be collected semiautomatically (**figure 2**).

Contact Frequency

Contact frequency is split into two categories:

1. **Past incident time series**—Consists of historical data from a security incident and event management (SIEM) solution or TIP, such as the number of incidents from the same advanced persistent threat (APT) group over a certain period of time.
2. **Past victims' geosectoral profile**—Takes into account several characteristics of previous victims of successful attacks. These metrics reveal whether attacks are hitting enterprises in similar sectors or are coming closer geographically. Typically, these data can be shared by using a TIP. The languages used in prior attacks may be relevant as well. APT groups might be successfully targeting enterprises through phishing emails in a certain language.

Figure 2—Threat Metrics by FAIR Category

FAIR Category	Evidence-Based Threat Category	Metric
Contact frequency	Past incident time series	Incident count within a period
		Trend change in the incident count
		Proportion of attacks by the actor over total attacks of the same type
		Trend change in the proportion of attacks by the actor over total attacks of the same type
	Past victims' geosectoral profile	Average match ratio of past victims' region within a period
		Trend change in the average match ratio of past victims' region within a period
		Average match ratio of past victims' country political alliance
		Average match ratio of past victims' country development level
		Average match ratio of past victims' country language
		Average match ratio of past victims' sector
		Trend change in the match ratio of past victims' sector
Probability of action	Threat actor's objective	Match ratio of an objective
		Trend change in the match ratio of the objective
	Threat actor's commitment	Days since the campaign started
		Days since the last attack
		Average number of days between attempts within a period
		Trend change in number of days between attempts within a period
Threat capability	Threat actor's skills	Sophistication level within a period
		Trend change in sophistication level
		ATT&CK coverage
		Trend change in ATT&CK coverage
		Efficiency
		Trend change in efficiency
Resistance strength	Detection capabilities	Campaign analysis: average kill chain detection phase
		Campaign synthesis: kill chain detection coverage
		DETT&CT overall coverage
		DETT&CT campaign coverage
	Exploitation surface	General exploitation surface
		Campaign exploitation surface
	Postdetection capabilities	Average investigation time
		Average response time

Probability of Action

The probability of action describes the likelihood that the threat actor will initiate an attack against a particular organization. The attractiveness thereof depends on the objective that the threat actor pursues in this particular campaign, the match of this objective to the organization (e.g., a hacktivist group might be more interested in attacking a pharmaceutical enterprise) and the threat actor's commitment to reach these objectives.

Threat Capability

Threat capability focuses on different characteristics of the threat actor behind the campaign. For instance, STIX defines seven sophistication levels, ranging from none to strategic.[7] In addition, the threat actor's known capabilities can be mapped to MITRE ATT&CK techniques, which are considered indicative of the APT group's capabilities.[8] A third component of threat capability is efficiency, which can be estimated by examining threat reports and security news to determine how many times the APT group carried out successful attacks.

Resistance Strength

Resistance strength is an organization-specific metric that defines the enterprise's defense capabilities. It is divided into three parts:

1. Detection capabilities are measured using the Lockheed Martin kill chain in two ways: first, by giving a higher score to an early detection in the kill chain, and then by scoring the overall coverage of the entire kill chain.[9] Alternatively, the enterprise's detection capabilities can be measured using the DETT&CT framework.[10] Because the threat actor's capabilities have been mapped to MITRE ATT&CK techniques, the enterprise's available defense techniques can be identified based on its DETT&CT capabilities.
2. Exploitation surface is divided into the general exploitation surface and the campaign exploitation surface. The latter is the exploitation surface mapped to the MITRE ATT&CK techniques used in the APT group's previous attacks.
3. Postdetection capabilities focus on the enterprise's security operations center (SOC) and computer security incident response team (CSIRT) capabilities. They take into account the average investigation time and the average response time for an incident.

Calculating an Aggregated Threat Score

Based on all these metrics, threats can be prioritized by calculating a single threat score per threat using Bayesian Belief Networks (BBNs). BBNs are a powerful knowledge representation and reasoning tool under conditions of uncertainty. A BBN is a directed acyclic graph (DAG) with a conditional probability distribution for each node.[11]

The FAIR taxonomy in **figure 3** is an example of a DAG. The metrics at the bottom of the graph are

Figure 3—Subset of the Metrics Implemented in the Bayesian Belief Network

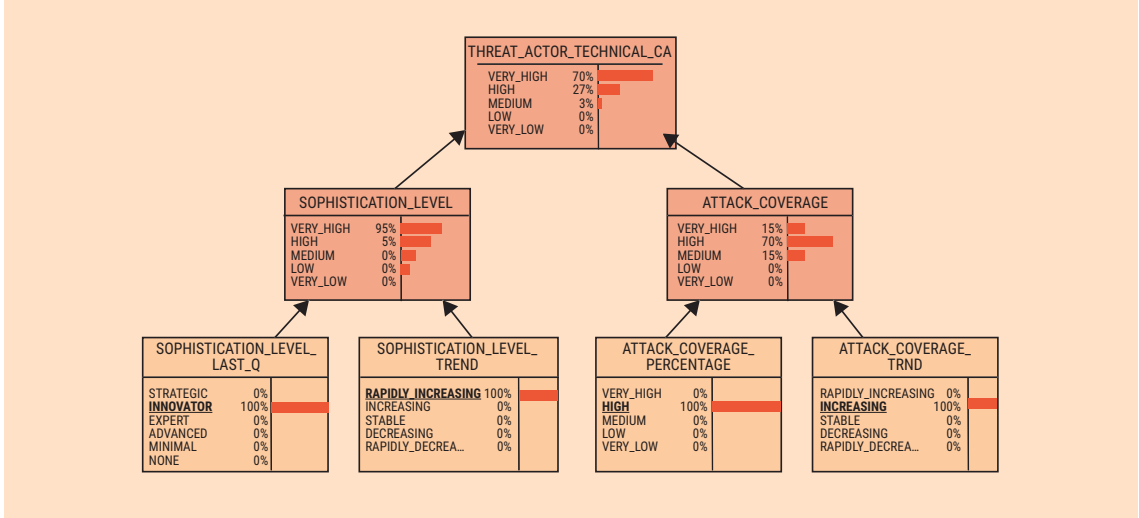

Enjoying this article?

- Read *Cyberrisk Quantification.* *www.isaca.org/cyberrisk-quantification*
- Learn more about, discuss and collaborate on information and cybersecurity in ISACA's Online Forums. *https://engage.isaca.org/onlineforums*

Emotet

Emotet's initial attack vector is delivered via infected email attachments through a fully automated process. During the lateral movement phase, it spreads across devices in the network, making it more resilient and, thus, more difficult to remove from the network. A plethora of servers around the world are used for various functions such as to spread the malware to new victims, act as command-and-control servers and make the malware more resistant to takedown attempts. The malware is polymorphic in nature, which makes it difficult to detect for signature-based defense mechanisms, since the malware changes its code each time it is called.[12, 13]

Emotet was originally designed in 2014 as a banking trojan intended to steal financial data, but it was later offered for hire to other cybercriminals, allowing them to install other types of malware, such as ransomware, onto a victim's computer.[14, 15]

LockerGoga

In the case of LockerGoga, the attackers seem to already know some of the targets' credentials at the start of an intrusion. These credentials may have been obtained through a successful phishing campaign or simply by buying them from other hackers.[16] The malware does not support any self-propagating code during the lateral movement phase.[17] LockerGoga partially encrypts files on the infected computer and leaves a ransom note on the user's desktop containing an email address, presumably so the victim can contact the attacker for decryption and payment options.[18]

The LockerGoga ransomware was first publicly reported in January 2019, when it was tied to an attack against the French engineering company Altran Technologies.[19] Other mutations were used in attacks against Norwegian aluminum manufacturer Norsk Hydro and two chemical companies, Hexion and Momentive.[20, 21]

connected to intermediate nodes by arcs, representing probabilistic dependencies. These probabilistic dependencies are expressed between two variables in the conditional probabilistic table. In this way, probabilities propagate to the top of the graph. For instance, at the top level of this BBN, the probability distribution of the threat score can be described as the relation between an expected loss magnitude and loss event frequency.

For this version of the threat prioritization methodology, a discrete BBN was used, which was based on categorical variables only (using labels such as "very low" and "high"). Although some data are lost in the transformation from continuous to categorical variables, it allows the analysis to be more enterprise specific and tailored to its environment. For instance, an incident count of 10,000 may be high for one enterprise but very low for another.

In the first manual round of exploration, a *pro forma* validation was performed using the evidence-based threat metrics combined with the BBN. For this exercise, two fictitious financial institutions (FIs) were described: one international bank and one smaller local bank. For both FIs, threat scores were calculated for two different families of malware actively used in campaigns by various threat actors: Emotet and LockerGoga. Emotet is a relatively advanced type of malware that can be configured as a banking trojan. LockerGoga is a far less advanced type of ransomware that does not specifically target the financial sector.

Based on the resulting threat scores in **figure 4**, the conclusion can be drawn that the LockerGoga malware is a bigger threat than Emotet for the local bank. For the international bank, Emotet is a greater threat than LockerGoga. The differences in outcomes are relatively small. This highlights the importance of using a methodology that can

Figure 4—Threat Scores

	LockerGoga	Emotet
Local bank	3.79	3.15
International bank	3.25	3.59

discern between threats in a reproducible way, based on metrics. It is impossible to prioritize threats correctly without proper observations, metrics and a methodology to rank threats.

Case Study: Sectoral Threat Landscape

The concept of metrics-driven threat scores was first applied in practice in the sectoral cyberthreat landscape of the Dutch finance industry. Through the so-called 1 Financial Threat Landscape for The Netherlands (1FTL-NL) initiative, leading FIs in the Netherlands jointly monitor the evolution of cybersecurity threats and the impact that developments in the threat landscape might have on their sector. The underlying purpose is to offer guidance for smaller FIs that might not have the means to maintain a self-reliant threat intelligence capability and to encourage collaboration on threats that affect the industry as a whole (e.g., by aligning individual intervention strategies). The 1FTL-NL threat landscape is compiled annually and subjected to an end-of-year review to identify lessons learned and improve future editions.

Similar to other initiatives, early editions of the 1FTL-NL threat landscape featured priority designations based on the expert insights of its formulators—a core group of specialists delegated by the participating FIs. Feedback from the 1FTL-NL target audience revealed a widespread desire to make the process of prioritization more transparent and less dependent on human opinions. In view of this, the 1FTL-NL initiative embraced the evidence-driven prioritization model. Rather than pursuing all 31 threat metrics in the model, 1FTL-NL focused on a selection of metrics (initially five, and later increased to eight) that could realistically be surveyed across its diverse constituency. The philosophy was to start relatively small and possibly refine the prioritization mechanism over time. To ensure a sufficiently balanced outcome, the selected metrics covered all the core categories of the prioritization model (**figure 2**), and the scale for each metric (i.e., the definition of "high," "medium" or "low") was tailored to the sectoral (rather than enterprise specific) nature of the 1FTL-NL landscape. **Figure 5** depicts the resulting BBN structure for aggregating the metrics' values into an overall sectoral threat score.

> "IT IS IMPOSSIBLE TO PRIORITIZE THREATS CORRECTLY WITHOUT PROPER OBSERVATIONS, METRICS AND A METHODOLOGY TO RANK THREATS."

Source input for the 1FTL-NL landscape was collected through written questionnaires sent to individual FIs and a selection of industry bodies. Respondents were asked to describe five "major threats" that they perceived as particularly relevant and then substantiate this selection by scoring the various threat metrics. To streamline this process, the questionnaires included intuitive descriptions of possible metric values. **Figure 6** shows an example of the format used.

Collaborative Cybersecurity Research With Dutch Industry

The work presented here stems from the Shared Research Program Cyber Security run by TNO (Netherlands Organisation for Applied Scientific Research) and the financial industry in the Netherlands between 2015 and 2020.[22] Within the context of this program, the model for the evidence-based prioritization of cybersecurity threats was compiled in close collaboration with cyberintelligence specialists at ABN AMRO, ING Bank NV, Rabobank and Volksbank. The 1FTL-NL was a separate initiative created to institutionalize a single, harmonized threat landscape for all FIs in the Netherlands. In view of their synergies, the two projects chose to collaborate and align where appropriate. TNO and the 1FTL-NL team are still collaborating to refine the model for metrics-driven threat priorities. The Shared Research Program was recently succeeded by the Partnership for Cyber Security Innovation (PCSI), which follows a novel process for staged innovation and features many of the same FIs.[23] The PCSI features most of the FIs that participated in the preceding SRP but intends to build an ecosystem for cybersecurity innovation in which partners from any industry are welcomed.

Figure 5—Concise BBN Structure for Calculating Sectoral Threat Score

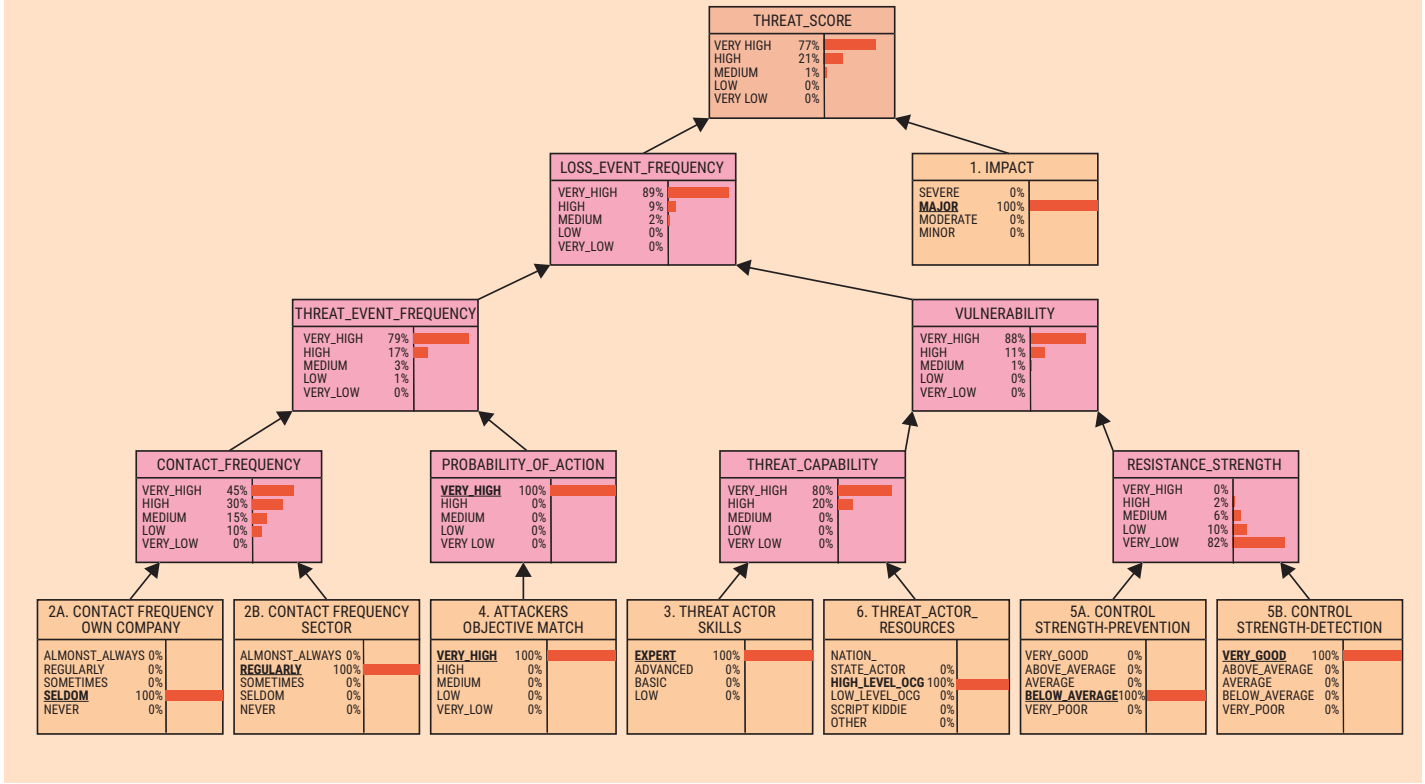

Figure 6—Sample Metric Employed in the 1FTL-NL Threat Landscape

2A. Contact Frequency Own Company	
How often has this particular threat materialized in your company? (Including unsuccessful attempts)	
NEVER	We have not seen any attempt thus far.
SELDOM	We saw a single attempt.
SOMETIMES	We saw a small number of attempts.
REGULARLY	We see this several times per month.
ALMOST ALWAYS	We see this on a daily basis.

Responses to the questionnaires were consolidated into a prioritized threat overview by the 1FTL-NL team. The consolidated values of individual threat metrics (e.g., the example shown in **figure 6**) were typically the result of expert interpretation rather than any mathematical operation. However, aggregated threats scores were compiled by feeding these values into the BBN calculation structure shown in **figure 5**.

The transition toward evidence (metrics)-based threat prioritization was perceived as a major step forward in the maturity of the 1FTL-NL initiative. Although it still involves some degree of expert judgment, the present prioritization scheme offers transparency in the ranking of threats. It also results in more consistency in the individual inputs supplied by 1FTL-NL constituents because they all assess threats based on the same factors and value scales. On the whole, the enhanced priority mechanism has increased both the credibility and the acceptance of 1FTL-NL as a leading source of threat insights for the Dutch finance industry.

Conclusion

It is both feasible and valuable to prioritize cybersecurity threats on the basis of evidence (observations) rather than human opinions. Metrics-driven threat priorities fulfill a widely perceived need for transparency, and they reveal nuances in the relative severity of threats that human experts might find hard to distinguish. Enterprises that

maintain sufficiently mature cybersecurity capabilities (typically those with reasonably well-established threat intelligence practices) should also find that the source information needed to quantify the respective threat metrics is realistically attainable.

Despite its promise, the presented model needs further validation and refinement before it can be adopted at scale. First and foremost, the process of obtaining the source information through which the 31 threat metrics can be valued must be automated; although manual retrieval is possible, it is very time-consuming. Pilot projects in actual operational environments might prove valuable to assess how (and to what degree) such automation can be implemented. In parallel, the metrics themselves might be refined. For example, the categorical (high-medium-low) metrics employed in the presented model might be replaced by a more continuous format and might even include "fuzzy" values when experts have conflicting evaluations.[24]

> "IT IS BOTH FEASIBLE AND VALUABLE TO PRIORITIZE CYBERSECURITY THREATS ON THE BASIS OF EVIDENCE (OBSERVATIONS) RATHER THAN HUMAN OPINIONS."

Acknowledgments

The authors would like to thank the 1FTL-NL core team for its kind contribution to the case study included in this article.

Endnotes

1 Schwartz, M. J.; "Emotet Malware Returns to Work After Holiday Break," *BankInfoSecurity*, 18 January 2019, *https://www.bankinfosecurity.com/emotet-malware-returns-to-work-after-holiday-break-a-11955*

2 Europol, "World's Most Dangerous Malware Emotet Disrupted Through Global Action," 27 January 2021, *https://www.europol.europa.eu/newsroom/news/world's-most-dangerous-malware-emotet-disrupted-through-global-action*

3 *Ibid.*

4 Malwarebytes, "Emotet," July 2020, *https://www.malwarebytes.com/emotet/*

5 O'Brien, J. D.; "Targeted Ransomware: An ISTR Special Report," Symantec, July 2019, *http://images.mktgassets.symantec.com/Web/Symantec/%7Bb464dc43-2ae0-4912-8758-b153d8f278e7%7D_Targeted_Ransomware_2019July.pdf*

6 *TrendMicro*, "What You Need to Know About the LockerGoga Ransomware," 20 March 2019, *https://www.trendmicro.com/vinfo/us/security/news/cyber-attacks/what-you-need-to-know-about-the-lockergoga-ransomware*

7 Harbison, M.; "Born This Way? Origins of LockerGoga," Unit 42 Blog, Palo Alto Networks, 26 March 2019, *https://unit42.paloaltonetworks.com/born-this-way-origins-of-lockergoga/*

8 Ilascu, I.; "New LockerGoga Ransomware Allegedly Used in Altran Attack," *BleepingComputer*, 30 January 2019, *https://www.bleepingcomputer.com/news/security/new-lockergoga-ransomware-allegedly-used-in-altran-attack/*

9 Beaumont, K.; "How LockerGoga Took Down Hydro—Ransomware Used in Targeted Attacks Aimed at Big Business," DoublePulsar, 21 March 2019, *https://doublepulsar.com/how-lockergoga-took-down-hydro-ransomware-used-in-targeted-attacks-aimed-at-big-business-c666551f5880*

10 Greenberg, A.; "A Guide to LockerGoga, the Ransomware Crippling Industrial Firms," *Wired*, 25 March 2019, *https://www.wired.com/story/lockergoga-ransomware-crippling-industrial-firms/*

11 TNO, Shared Research Program (SRP) Cyber Security, *https://www.tno.nl/en/collaboration/partners-of-tno/shared-research-program-srp-cyber-security/*

12 European Network and Information Security Agency (ENISA), "ENISA Threat Landscape 2020," April 2020, *https://www.enisa.europa.eu/topics/threat-risk-management/threats-and-trends*

13 Global System for Mobile Communications (GSMA), "Mobile Telecommunications Security Threat Landscape," January 2020, *https://www.gsma.com/security/wp-content/uploads/2019/03/GSMASecurity-Threat-Landscape-31.1.19.pdf*

14 Freund, J.; *Measuring and Managing Information Risk: A FAIR Approach*, Butterworth-Heinemann, United Kingdom, 2014
15 Oasis Open, "Committee Specification 02," STIX™ Version 2.1, 25 January 2021, *https://docs.oasis-open.org/cti/stix/v2.1/stix-v2.1.html*
16 *Op cit* Freund
17 Open Group Library, "Risk Taxonomy (O-RT), Version 2.0," 18 October 2013, *https://publications.opengroup.org/c13k*
18 *Op cit* Oasis Open
19 MITRE Corporation, "Enterprise Techniques," *https://attack.mitre.org/techniques/enterprise/*
20 Hutchins, E. A.; "Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains," *Leading Issues in Information Warfare and Security Research*, January 2011, *https://www.lockheedmartin.com/content/dam/lockheed-martin/rms/documents/cyber/LM-White-Paper-Intel-Driven-Defense.pdf*
21 Marcus Bakker, R. B.; "Rabobank-cdc/DeTTECT," Github, March 2021, *https://github.com/rabobank-cdc/DeTTECT*
22 Le, A. C.; "Incorporating FAIR Into Bayesian Network for Numerical Assessment of Loss Event Frequencies of Smart Grid Cyber Threats," *Mobile Networks and Applications*, vol. 24, iss. 5, 2019, p. 1713–1721, *https://link.springer.com/article/10.1007/s11036-018-1047-6*
23 *Ibid.*
24 Partnership for Cyber Security Innovation, *https://pcsi.nl/*