IMPROVING THE SPEED OF CYBER DEFENSE:

SEVEN CYBER METRICS EVERY OWNER-OPERATOR NEEDS TO KNOW

CONTENTS

You are a bridge officer on a 250 m (820 ft), 12,000 TEU container vessel in the Kill Van Kull Channel, heading west towards the Bayonne Bridge. It is a cloudless, sunny day and you are cruising at 4 knots in your traffic lane, with a pilot on board. As you pass the Constable Hook Reach, you feel the ship start to veer hard to port and your speed appears to increase. Although the bridge instruments show your expected speed, location, and rudder position, the rudder is, in fact, hard over to port and your speed has increased to 12 knots. The pilot's PPU shows the vessel horribly deviating from the assigned course and speed, adding to the confusion. You immediately pull back on the throttle and turn the helm to hard right rudder, but the ship does not respond. Within a few minutes, your bow has run aground on the south shore of the channel, while the stern continues to swing around towards the north shore. Within six minutes, your ship is sideways in the channel and traffic in both directions has come to a halt.

So, what happened here? System malfunction? Crew failure? Cyberattack? How would you know and how would you tell the difference?

This scenario—inspired by our colleagues at the Cyber-SHIP Lab at the University of Plymouth, UK—is, actually, a potential cyberattack. It is not one of the doomsday attacks you see in the action movies, where an international criminal mastermind somehow takes remote control of your ship in an attempt to ransom the owners. It's actually much worse because it's more plausible. Malware on a shipboard system could actualize the very events described here... not only at the Ports of New York and New Jersey, but at any soft maritime target that a malign actor chooses. There are many ways that such malware can infect ships, ranging from spearphishing emails with bogus chart updates to device firmware patches from compromised vendors. And, like the Solar Winds attack demonstrated two years ago, multiple ships can be targeted at the same

time with malware that lies dormant for weeks or months.

When trying to manage cyber threats and understand cyber attack capabilities, maritime executives will benefit by measuring their respective risks and opportunities in temporal terms. In the commercial sectors with the greatest degree of cybersecurity maturity, corporate information security executives typically reach for their stopwatches to measure how robust their cybersecurity posture is. It is past time that maritime executives start doing the same. Here are seven cyber metrics that maritime owners and operators should understand and start to measure.

BREAKOUT TIME

Cyber "breakout time" was first defined by the cybersecurity firm Crowdstrike as "the time taken by an adversary to move laterally, from an initially compromised host to (an)other host(s) within the victim environment."[1] To use a modern piracy analogy, this is the time it takes for the pirates to get to the pilot house and take control of the ship starting from when they first anchor their grappling hook to the bulwark. As a matter of threat intelligence, this measurement is critical to understanding the danger that particular threat actors pose to platforms, businesses, and operations alike. According to Crowdstrike, in 2018 the fastest average breakout time was demonstrated by Russian state actors who were able to conduct full cyber intrusions (gaining initial access, escalating privileges, and gaining control over key systems and information on a target network) in 00:18:49 (hours:minutes:seconds). This was followed by the second-fastest, North Korean state actors, at 02:20:14, and the third-fastest, Chinese-state actors, at 04:00:26[2]. As adversaries advance their tools and tradecraft, breakout times continue to shrink, requiring defensive teams to adopt more automation and 24/7 operational postures.

MEAN TIME TO DETECT

Mean Time to Detect (MTTD) is breakout time's defensive counterpart. It measures the time network defenders require to identify a cyber breach into a defended network. If an adversary's breakout time is faster than a defender's MTTD, it means that the adversary is allowed unobserved, undefended actions against a platform or network for the length of time between the two measures. If MTTD is less than an adversary's breakout time, then attempts at further compromising high-value network assets and systems, data destruction, or harm to physical systems might be stopped (or impacts minimized) by the defenders. To be clear, detection does not mean that an attack is defeated; detection merely means that defenders recognize that a breach or attack is occurring and allows for the possibility of meaningful response.

In other sectors, MTTD is reduced by the around-the-clock management of network sensor data by a security operations center (SOC), which continuously review log data and network performance information to identify anomalies that indicate the early stages of a compromise. For most commercial fleets, this degree of continuous, real-time monitoring is not yet viable. More viable for most—but still daunting for many—is the regular evaluation of firewall logs, key host logs, and logging data from a vessel's highest-value systems at the pace at which such logs can be passed from vessel-to-shore. Too often, MTTDs are unnecessarily lengthened when organizations do not have processes for continuous log analysis on otherwise available shipboard data, allowing malicious cyber actors additional time to exploit a successful access attempt into something more costly and disruptive.

MEAN TIME TO PATCH

Mean Time to Patch is the measure of how long it takes, on average, to update systems with patches that have been publicly released. This measurement is threat agnostic and primarily a measure of an organization's baseline defensive cyber agility. While sophisticated or well-resourced actors can find novel vulnerabilities in systems for which patches do not yet exist (typically referred to as "zero day vulnerabilities"), threat actors of all levels of sophistication exploit the lag between the announcement of a patched vulnerability and application of the patch by defenders.

Analysis published by Mandiant in 2020 demonstrated that 42% of vulnerabilities are still exploited well after patches are publicly released (Fig 1), with the vast majority of observed cyber intrusions leveraging vulnerabilities for which patches are available but unapplied (Fig 2)[3].

[1] From Crowdstrike: "Lateral movement refers to the techniques that a cyber attacker uses, after gaining initial access, to move deeper into a network in search of sensitive data and other high-value assets. After entering the network, the attacker maintains ongoing access by moving through the compromised environment and obtaining increased privileges using various tools. Lateral movement allows a threat actor to avoid detection and retain access, even if discovered on the machine that was first infected. And with a protracted dwell time, data theft might not occur until weeks or even months after the original breach." https://www.crowdstrike.com/blog/the-myth-of-part-time-threat-hunting-part-1/

[2] https://www.crowdstrike.com/blog/first-ever-adversary-ranking-in-2019-global-threat-report-highlights-the-importance-of-speed/

[3] https://www.mandiant.com/resources/time-between-disclosure-patch-release-and-vulnerability-exploitation

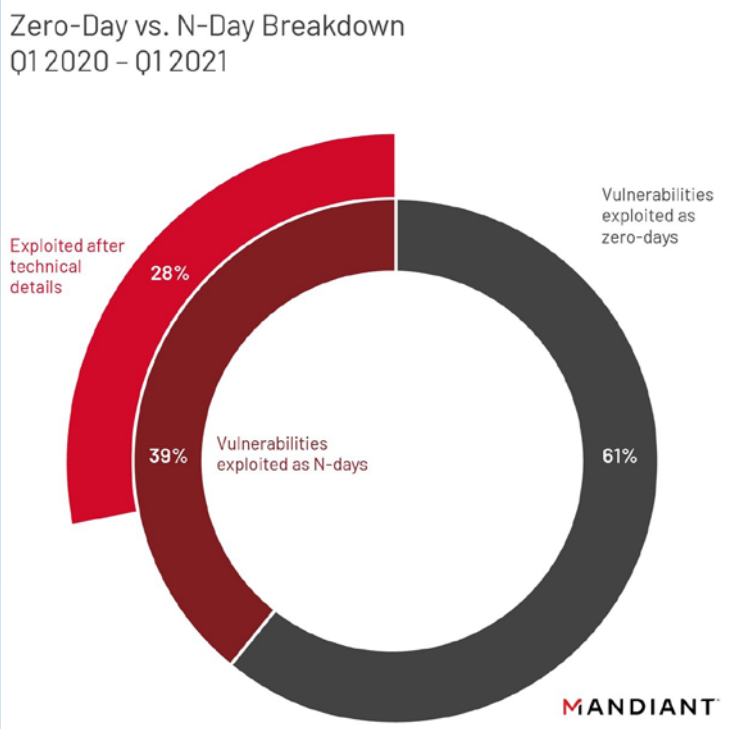

Figure 1: The cybersecurity firm Mandiant analyzed 95 vulnerabilities that were exploited between Q1 2020 through Q1 2021 and found that a majority (61 percent) of analyzed exploited vulnerabilities were exploited as zero-days, and more zero-days were exploited in the first quarter of 2021 than each of the past four years. Forty-one (41) percent of n-day vulnerabilities (those exploited for the first time after patch publication) were first known to be exploited within three weeks of patch availability.

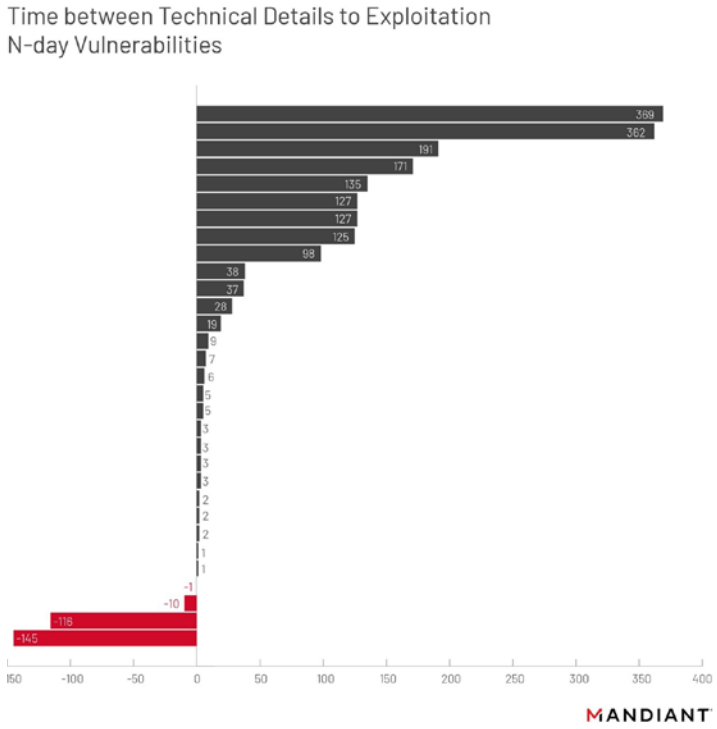

Figure 2: A more granular examination over the same activity monitored by Mandiant depicted in Fig 1. Many patches are released just days after or at the same time vulnerabilities are made public, which explains the clustering of exploits being used just before or right around the time a patch is released. Patches remaining unapplied years after being released is not uncommon, allowing low-sophistication cyber actors the ability to successfully attack, absent having to invest in vulnerability discovery programs or purchase "zero-day" vulnerabilities from other hacking groups.

Masters and vessel owners oftentimes rely on off-ship service providers to protect their vessel before malicious network traffic reaches the ship or at firewalls that exist at the internet-boundary of shipboard networks. While this is not generally the same type of cyber defense as patching vulnerable systems on a vessel, the temporal metrics are similar. Most off-ship defenses will rely on proprietary vulnerability and threat information that allows defensive rule-sets to be put into place that block incoming messages to ships that fit known-malicious profiles. The amount of time it takes for a vulnerability or a specific adversary tactic to be initially identified, reported, translated into defensive-rule sets, and then put into place follows the same dynamics of MTTP. Adversaries that only use known vulnerabilities can still be faster than off-ship cybersecurity service providers.

MTTP, however, is not merely a matter of availability of patches and personnel, but can be a impacted by the will of the organization and other business imperatives. The USCG 2021 Cyber Trends in Maritime Report[4] found that MTS partners fully remediated 66 per cent of all publicly exploitable findings and 45 per cent of internally exploitable findings within six months of detection. But that means that one-third of publicly exploitable findings and more than half of internally exploitable findings were not fully mitigated a half-year later.

MEAN TIME TO CONTAIN

These measures define the time it takes from detection until either on-site or remote defensive forces can respond to the detection of malicious activity. Mean Time to Contain (MTTC) describes how long it takes the response effort to contain the attacker's attempts at conducting further actions or exploiting additional systems. This typically means that defenders have successfully identified the root-cause of a compromise and mitigated both the adversaries' actions on their network as well as the underlying vulnerability(-ies) that enabled the compromise in the first place. Combined with mean time to detect, the time to contain a threat effectively describes the amount of time a notional

[4] https://www.dco.uscg.mil/Portals/9/2021CyberTrendsInsightsMarineEnvironmentReport.pdf

adversary has until a given attack can be stopped.

Alternatively, MTTC is the amount of time it takes between detection of an event and the determination that whatever indicator(s) initially tipped off defenders is actually a false-positive. According to BitDefender, approximately 50% of cybersecurity alerts in organizations with full time security operation centers are actually false positives[5]. As a result, larger organizations may spend as many as 400-hours per week resolving events as false positives. While well-crafted alert rules and expertly tuned sensors will reduce false positives, the dynamic nature of contemporary maritime systems and networks means that false positives are all but impossible to eliminate. Despite shipboard traffic being far less complex and smaller in volume than the corporate networks described in the aforementioned studies, being able to identify and quickly resolve false positive events is as critical as the speeds associated with responses to actual security events.

MEAN TIME TO RECOVER

Mean Time to Recover (MTTR) is the average time it takes for an organization to recover from a successful cyber attack. A ship whose computers and control systems are completely crippled by ransomware, for example, will require all the affected systems to be wiped, restored, replaced, or repaired as well as the assistance of cybersecurity professionals to identify and fix the root cause(s) of the ransomware infestation.

While recovery from cyber attacks varies based on the nature of the affected organization (e.g. geographically-dispersed organizations often require teams to physically travel to each affected site for recovery) and the nature of the disruption, recovery timelines currently tend to be measured in terms of weeks-to-months for most commercial and industrial organizations. Maersk, for example, took nine days to partially recover from the NotPetya attack in 2017 and full recovery took months. Speaking two years after the attack, Maersk's Chief Information Security Officer stated that the company was now attempting to build the internal capacity to recover from a similar attack in 24 hours.[6]

The problem is exacerbated on a vessel since the computers are often specialized and/or embedded systems. Whereas an infected computer in a typical office can be quickly replaced and get employees back working, imagine the time to recover from an attack on a specialized ship system without backups onboard or even back at home port.

HOW FAST IS "FAST ENOUGH?"

CrowdStrike currently defines companies as "top performing" cyber organizations if they can meet the "1-10-60 Rule": detecting an intrusion within 1 minute, investigating within 10 minutes, and containing or remediating the problem within 60 minutes. When adversaries are allowed to engage in unchecked lateral movement over a protracted period, the likelihood of adversary success goes up significantly.[7] CPO Magazine, citing statistics from Mandiant, notes that most organizations are not anywhere near "top performing." "The average containment time in 2017 was five days for network intrusions. However, that is the containment time after detection. The average time to detection was a staggering 66 days."[8]

Otherwise stated, the average organization could be compromised for over 5000 times longer than it would take a Russian cyber team to complete their objectives before a defensive cyber response would even begin.

In non-maritime commercial enterprises, "acceptable" cyber metrics are determined by the business and/or sector-specific requirements governing the fiscal and commercial viability of that company. While an airline might be able to recover from a days-long denial-of-service attack on its ticketing system, for example, a month-long outage would likely collapse most commercial carriers. Thus, acceptable MTTC/MTTR are determined by the fiscal realities of the business, not cyber capabilities.

For commercial fleets, determining acceptable MTTC/MTTR largely follows the economic contours of other commercial enterprises, with two additional maritime-specific factors. First, any inability to get assets into port at scale takes on an immediate fiscal urgency for shipping fleets. Infamously, Hanjin Shipping only lasted four days between receivership preventing Hanjin ships from entering ports and its bankruptcy[9]. Similarly, a fleet-wide cyber disruption to shipping operations could result in potentially existential fiscal impacts to commercial fleets in a matter of days, with crippling second-order global economic effects. Second, cyber disruptions in restricted maneuvering situations like the hypothetical one at the beginning of this article can be catastrophic. A cyber-physical disruption that

[5] https://businessinsights.bitdefender.com/every-hour-socs-run-15-minutes-are-wasted-on-false-positives
[6] https://www.csoonline.com/article/3444620/rebuilding-after-notpetya-how-maersk-moved-forward.html
[7] https://www.crowdstrike.com/cybersecurity-101/lateral-movement/
[8] https://www.cpomagazine.com/cyber-security/hacks-are-happening-faster-how-much-do-cyber-response-times-need-to-improve/
[9] https://www.seatrade-maritime.com/americas/hanjin-shipping-files-us-bankruptcy-protection

prevents normal maneuvering of a ship quickly becomes a catastrophic matter when paired with extreme weather or unfavorable wind and currents in a restricted passageway. The Ever Given grounding in the Suez Canal, for example, was a result of high winds and what was reportedly a temporary loss of steering control lasting only seconds, all occurring in the middle of the day with two pilots aboard.

Masters, fleet operators, and maritime executives from the unit level and up should be aware of their respective MTTD, MTTP, MTTC/MTTR, even if dependent on outside service providers for one or all the aforementioned defensive capabilities. Absent on-ship or secure remote capabilities to recover lost systems or networks following an attack, a master should be knowledgeable of their vessel's likely MTTRs timelines if for no other reason than to be better able to manage the subsequent risks to their ship and crew through the duration of the degradation.

At the fleet operator level–and assuming a limited capacity for cyber response and recovery–shipping line staff and operators should understand how cyber impacts against multiple vessels will affect MTTC/MTTR timelines. If these timelines are going to be dependent on vessels reaching ports where technicians can then assist the vessel, then it may be prudent to ensure that divert and response plans are coordinated such that technicians are able to meet vessels within their typical laytimes.

CYBER GEOGRAPHY

Even smaller, family-owned fleets operate globally, therefore there is the likely further challenge of responding to geographically dispersed cyber incidents. A single cyber response team cannot respond to cyber-induced shipboard casualties in both the Atlantic and Indian Oceans at the same time, even if they are caused by the same underlying system vulnerabilities and/or adversary actions. Consequently, the number and forward pre-positioning of cyber-response capabilities required to respond to geographically dispersed cyber casualties requires careful consideration.

As such, a foundational planning concern for 21st century shipping lines is establishing and preserving the ability to scale and mobilize the cyber incident containment and recovery to affected platforms inside of acceptable timelines. At present, however, it does not appear that it is possible to train, hire, and retain enough cybersecurity professionals such that an expert cyber team can be assigned as ship's company on every platform. The more viable construct is the rapid mobilization of cyber response and recovery teams operating from a central hub close enough to a lines' key shipping routes to enable teams to be 'waiting on the pier' to support affected vessels as soon as they reach their next port of call.

CYBER WORKFORCE HEALTH: THE MOST IMPORTANT METRIC OF ALL

This article has frequently referenced "defenders" and "response and recovery teams," despite most maritime owner-operator entities lacking any type of formal cybersecurity teams. The efficiency and speed of such response teams can be the difference between nuisance and disaster. Many commercial maritime entities are still early enough in their respective cybersecurity journeys to only now consider hiring their first full-time cybersecurity professionals. Often these are management-level personnel such as Chief Information Security Officers (CISOs) tasked with managing company-wide cybersecurity risk, as opposed to teams of personnel prepared to respond to security incidents.

Even with a cybersecurity manager in place, owner-operators will lack sufficient capacity to track and respond to cybersecurity events and incidents. Absent creating a team of cybersecurity professionals to conduct full-time security operations on shipboard and shore-based networks, many maritime companies will likely find it prudent to look within their existing engineering and operational organizations to find groups of individuals interested in taking on cybersecurity responsibilities in exchange for training and education opportunities. Sometimes referred to as "re-skilling" or "up-skilling," taking employees already familiar with an organization's equipment, business practices, and culture and providing them with opportunities to develop cybersecurity expertise has significant advantages when compared to attempting to hire outside professionals into a maritime organization. Globally, as of 2021 there was a cybersecurity workforce shortage of over 2.72 million professionals, with the shortfall increasing year after year for more than a decade[10]. Finding quality, affordable cyber professionals is out of reach for most companies, say nothing of those whose employees work in industrial environments, are frequently required to travel, and potentially spend weeks out at sea. In addition, the cybersecurity workforce shortage describes cyberdefense in "generic" industries; there is a real dearth of professional who understand both cybersecurity and maritime, and this is not (yet?) a subject taught sufficiently at maritime academies. As an industry, we are forced to grow our own.

[10] https://www.nist.gov/system/files/documents/2022/07/06/NICE%20FactSheet_Workforce%20Demand_Final_20211202.pdf

For most maritime companies, first identifying those capital assets (e.g. key shipboard machinery and systems, key port infrastructure, etc.) that are reliant on computers, control systems, and/or internet connectivity for their proper function and most critical to business operations is the first step in identifying which portions of their respective workforce should be first considered for up-skilling opportunities. Once a company's most critical systems are identified, look at the pool of personnel who interact with these systems most frequently to try and identify any individuals who might be willing and able to receive tailored cybersecurity training. Once identified, send these individuals to any of the many cybersecurity training events that are offered regularly around the world. Cybersecurity training and education has developed a range of formal certifications and training progressions. Most of these certifications and trainings are offered via week-long seminars adjacent to conferences or other events. Generally, these are held at pleasant-to-visit, often family friendly locations that serve as built-in incentives for continued participation and skill progression.

While the timing and tempo of up-skilling opportunities will be resource and situationally dependent, a general goal for a maritime company is to attempt to grow at least three trained cyber professionals within each of the business divisions responsible for their highest-value assets. While members of other business divisions can aid in response and recovery, those upskilled individuals associated with a particular system or network will be the core team members to deal with emergent cyber incidents. Workforce sufficiency, therefore, is measured as a factor of the number of high-value, high impact systems and networks within an organization and the number of cybersecurity upskilled personnel familiar/expert on those systems.

SAFE WATER

You are a bridge officer on a 250 m (820 ft), 12,000 TEU container vessel in the Kill Van Kull Channel, heading west towards the Bayonne Bridge. It is a cloudless, sunny day and you are cruising at 10 knots in your traffic lane. A few minutes from Constable Hook Reach, you get a call from the Radio officer.

"Sorry to do this to you again, but I just got a call from corporate to block all traffic from the Business VLAN to the engineering and navigation VLANs. You won't have internet access from the Bridge laptop for the next ten minutes or so until I can put these new firewall rules in place. Just give me a call if you notice anything weird happen." They hang up.

You walk out to the bridge wing. You pass Constable Hook Reach at your expected cruising speed, staying precisely in the center of your traffic lane all the way to the terminal for a safe, on-time arrival.

About the Authors

Tyson B. Meadors is the Practice Lead for Maritime Cyber at Fathom5. Heralded as "America's Top Cybersecurity Strategist" by former National Security Advisor H.R. McMaster, he served as Director for Cyber Policy on the National Security Council Staff from 2017-2018, advising the President, Vice President, and multiple National Security Advisors and was the lead author of the 2018 U.S. Cybersecurity Strategy and key contributor to a range of Executive Branch strategies and policies. He is certified and experienced practitioner in a range of cybersecurity disciplines, to include incident response, forensics, penetration testing, vulnerability analysis, cyber-physical security engineering, assessment, and auditing. He also currently serves on the Editorial Board of the US Naval Institute. He has degrees and certifications from a range of institutions, to include the U.S. Naval Academy, the North China Institute of Science and Technology (华北科技学院), the US Naval War College, Old Dominion University, and the Escal Institute of Advanced Technologies (SANS).

Gary C. Kessler, Ph.D., CISSP, is a Principal Consultant in the Maritime Solutions Group at Fathom5. Co-author of Maritime Cybersecurity: A Guide for Leaders and Managers (2/e) and a retired professor of cybersecurity, he is the President of Gary Kessler Associates, a consulting, research, and training company located in Ormond Beach, Florida, a non-resident senior fellow at the Atlantic Council, a visiting faculty member at the U.S. Coast Guard Academy, and active in the USCG Auxiliary and AUXCYBER program. Gary holds a merchant mariner credential and is a SCUBA instructor.

If you are interested in measuring or enhancing your organization's cyber capabilities, please reach out to us at:
cyber@fathom5.co


FATHOM5
BRILLIANT MACHINES