

TAG Cyber

2022

Security Annual

SPECIAL REPRINT EDITION

WHY ISLAND BUILT THE ENTERPRISE BROWSER

AN INTERVIEW WITH MIKE FEY,
CO-FOUNDER & CEO, ISLAND

CYBERSECURITY METRICS: WHAT GOOD LOOKS LIKE

USING TIME PATTERNS TO PREDICT
FUTURE CYBERCAMPAIGNS

TAGCYBER
DISTINGUISHED



Island

The need to reduce cyber risk has never been greater, and Island has demonstrated excellence in this regard. The TAG Cyber analysts have selected Island as a 2022 Distinguished Vendor, and such award is based on merit. Enterprise teams using Island's platform will experience world-class risk reduction—and nothing is more important in enterprise security today.



The Editors,
TAG Cyber Security Annual
www.tag-cyber.com

**WHY ISLAND BUILT THE
ENTERPRISE BROWSER**

AN INTERVIEW WITH MIKE FEY,
CO-FOUNDER & CEO, ISLAND

3

CYBERSECURITY METRICS: WHAT GOOD LOOKS LIKE

Dr. Jennifer Bayuk

6

**USING TIME PATTERNS TO PREDICT
FUTURE CYBERCAMPAIGNS**

Dr. Edward Amoroso

11

REPRINTED FROM THE TAG CYBER SECURITY ANNUAL

©TAG CYBER 2023



AN INTERVIEW WITH MIKE FEY,
CO-FOUNDER & CEO, ISLAND

WHY ISLAND BUILT THE ENTERPRISE BROWSER

If asked to list their most critical applications, it's curious that many enterprise teams would forget to mention their browser. Perhaps because the browser is so obviously present in every environment, it is often taken for granted by security teams, compliance managers and requirements framework curators (like NIST).

Cybersecurity start-up Island provides an enterprise-grade browser that includes many valuable security features. The company focuses on so-called last mile protections which complement—or even replace—some existing endpoint controls. We spent time with the Island team to learn more about these exciting advances.


TAG Cyber: Why do you think so many companies take their browser for granted in the context of their security architecture?

ISLAND: It's not so much that they take the browser for granted. They take the browser very seriously, but the status quo for decades now is that they don't have a lot of control over the browser itself. Think about the browser compared to other domains in IT. We have so much control over the operating system, with the ability to configure and manage the OS to satisfy every enterprise requirement, but, by and large, the browser has not kept up, even though the browser is now running our most critical enterprise applications. This gap forced security teams to implement a whole host of security tools outside the browser—everything from web filters and DLP to virtualization and even remote browser infrastructure. Island introduced a completely different approach: We're building a browser that delivers those critical enterprise controls natively, inside the browser. We're giving the browser an active role in enterprise security.

TAG Cyber: What's the difference between an Enterprise Browser and a consumer browser?

ISLAND: An Enterprise Browser is built to integrate and cooperate with the enterprise. This approach delivers significant improvements to the security posture both by reducing complexity and increasing effectiveness. It also has a dramatic impact on IT organizations by playing a key role in delivering applications and resources to their users. Finally, it improves the experience for end-users with a consistent, fluid user experience and strong productivity enhancements. When it comes to the specific capabilities of the Enterprise

**We're working
in some of the
most challenging
enterprise
environments to
help our customers
make BYOD
a success.**



Browser, it's really about last mile controls—the ability to govern everything that happens at the presentation layer of the browser with dexterity and logic. Everything—from what a user sees to how they interact with applications and data—is now controlled by the enterprise, in ways that were never possible before.

TAG Cyber: How does your browser fit in with existing applications, websites and workflows?

ISLAND: The beautiful thing is that because we're built on the open-source Chromium browser engine, when you first engage with the Enterprise Browser, you'll feel no difference. It delivers the same user experience, look and feel, as well as 100% web app compatibility. There's no need for end-user training or documentation, because everyone already knows how to use a web browser. When you start to look at protecting data, enhancing user productivity and giving insights to security and IT organizations, that's where you start to see the differentiation. We've added a control and governance layer inside the browser to support any business objective. In this way, the end-users get a tool that feels very familiar to what they're used to, and the supporting functions gain a whole range of new capabilities.

TAG Cyber: You've now deployed in some well-known organizations, tell us how your customers have realized value from the Enterprise Browser?

ISLAND: Our customers have realized value in several different ways. For some, the value comes from an improvement in their security posture, like the customer who uses Island to satisfy several key objectives for their HITRUST requirements. For others, the value comes from dramatically simplifying their security stack, as well as reducing expenses and operational complexity. We've helped customers rethink the architecture for key business processes to take out layers of complexity and improve the experience for employees and customers. What I think is most exciting is that it's helping organizations embrace the future. It's helping them embrace BYOD for employee flexibility. It's helping them work with contractors and business process outsourcers (BPOs) in a new, more efficient model. Moreover, it's helping them think about SaaS apps as a safe place for business-critical data, knowing that they have full control over how and where data can enter or leave. While there's huge value and ROI in the reduction of complexity and improvement of security posture, what I think most organizations are seeing is an exciting new platform that allows them to embrace the modern workforce.

TAG Cyber: Can you share some insights into how the Enterprise Browser changes the security and IT landscape?

ISLAND: Right now, we're working in some of the most challenging enterprise environments to help our customers make BYOD a success. That's interesting for two reasons. First, BYOD is not a new idea. These customers have looked at many other approaches over the years and found them all lacking. Second, these are large customers with global footprints. If the Enterprise Browser can meet the challenges of the most difficult to support environment—the highly variable and distributed BYOD environment—then it can operate everywhere. If we can deliver sensitive data and business-critical applications in the most challenging BYOD model, and do it cost-effectively, how does that not disrupt the old paradigm of heavy, complex security stacks?

As we think about the next step on this journey, BYOD leads to "Self IT." As more digital-native workers enter the workforce, we can expect more employees to manage their own IT to meet their particular needs. This is an opportunity to re-think our role in supporting the workforce. To put it another way, the Enterprise Browser completes the journey of SaaS. In the first wave of SaaS, we moved our applications out of data centers, and we stopped installing thick apps on the desktop. The role of IT operations shifted from managing physical servers to configuring SaaS resources. We have the opportunity to do the same thing with the endpoint: We don't need to physically handle each endpoint; we can instead configure and manage the operating system. In the SaaS model, the OS is not Windows, not macOS, not iOS. It's the browser. That's where work gets done.

CYBERSECURITY METRICS: WHAT GOOD LOOKS LIKE

DR. JENNIFER BAYUK

Measurement is the process of mapping from the empirical world to the formal, relational world. The measure that results characterizes an attribute of some object under scrutiny. Cybersecurity is not the object of measurement, nor a well-understood attribute. This means you are not directly measuring security, you are measuring other things and using them to draw conclusions about cybersecurity. Cybersecurity metrics can create situational awareness on multiple fronts. Some of it will be good news, some of it bad news. But note the difference between good *news* and good *metrics*. Bad metrics can be good news. In security we call that a “false negative.”

Good metrics can give both good and bad news that can be trusted. By good metrics, we mean metrics that are both practical and useful. We can learn from them and use them to systematically improve practices. Practical and useful metrics are easy to connect to the concept of cybersecurity. They utilize transparent data-gathering processes and support security decision-making.

Good cybersecurity (as opposed to good metrics) looks like swift and thorough cyberattack containment, mitigation and root-cause remediation. In the absence of attacks, good cybersecurity looks like a low risk of successful attack. A demonstration that attack response is good requires a *performance* metric. A conclusion that there is a low risk of attack requires a *goal* metric; that is, we operate under the assumption that the goal of a cybersecurity program is to reduce the risk of a negatively impacting cyber event to an acceptable level.

Sometimes this distinction between performance and goal metrics is described as “correctness versus effectiveness” or “verification versus validation.” Performance, correctness and verification measures are grounded in specifications for system composition. Goal, effectiveness and validation measures target whether the system accomplishes its mission. In systems engineering terms, performance metrics answer the question: “Was the system built right?” Goal metrics answer the question: “Was the right system built?”

In systems engineering terms, performance metrics answer the question: “Was the system built right?” Goal metrics answer the question: “Was the right system built?”



Industrial engineers intuitively understand that business-critical processes must be instrumented for measurement in order to be successfully managed. That's why pressure gauges on tanks used to measure capacity are typically customized and delivered with the tank rather than bolted on after the tank is integrated into its target environment. These measures, in combination with tank inventory, can show that the system is working as designed. Similarly, cybersecurity units of measure are tangible attributes of the cybersecurity ecosystem, such as information classification (nominal, a label), vulnerability exposure (ordinal, e.g. high, medium, low), server counts (numeric) or a time to respond to an incident (interval).

I reserve the term "measure" for acts of cybersecurity attribute data collection. When measures are combined via algorithms, metrics may be produced. Most performance metrics will use multiple measures. Figure 1 provides an example of using an algorithm to combine information classification, vulnerability exposure and server counts to calculate the percentage of servers with sensitive information that have critical vulnerabilities. Such measures of the control environment allow you to create algorithms that produces information you can use to see if your security program is operating as expected—that is, a verification (or lack thereof).




Figure 1: Measures + Algorithm = Metric

However, verification metrics do not convey information about risk. A percentage can be a risk measure only if it provides information about the probability that a system will succumb to attack. For that, you need information about threats as well as controls. Both percentages are ratios in that they have at least two measures: a numerator and a denominator. But only if the numerator feasibly approximates the chance of succumbing to attack at any given moment can the metric approximate risk. That is why so many publications and systems use the term "risk indicator" as opposed to "risk metric." The best performance indicator can only reflect whether the security was correctly built, not that it was adequate to thwart threats.

Cybersecurity practitioners often ignore this distinction and focus directly on finding and fixing security attributes that make them vulnerable, like common vulnerabilities. This focus results in metrics that look like Figure 2. Gary McGraw coined the term "Badness-Ometer" for this type of metric. It can only display poor security, never excellent security. The graph on the right of Figure 2 counts as a verification metric because it relies on counting vulnerabilities (bad things) in combination with a measure of time since the vulnerability was identified, and a time threshold set by management on how soon vulnerabilities

should be fixed. The three measures taken at monthly intervals add up to one metric that shows what bad looks like: the security performance target was not achieved. In the performance versus goal metric context, it shows that the system was not built right. There are also examples of Badness-Ometers that are goal metrics, such as annual monetary losses due to cyberattack (assuming your goal is not to have any).




Figure 2: McGraw's Badness-Ometer and an Exemplar Metric

Another readily accessible security attribute is found in security software obtained to meet a goal of system protection. But even this must be instrumented properly to produce a reliable performance measure. For example, it is common to set up a standard server-build process wherein security software such as antivirus or OS hardening agents are installed as part of a workflow. Successful completion of this step for all new and upgraded servers is often taken as a positive performance measure. It is also common for legacy machines to avoid this workflow by never upgrading or receiving the installation even though the security software is not able to run on the legacy OS. This leaves a pool of vulnerable servers below the radar of the measure. Only by careful enumeration of servers within scope and sufficient instrumentation on all servers to show what software is currently operational can you rely on performance measures to show what good performance looks like. Figure 3 illustrates the approach.

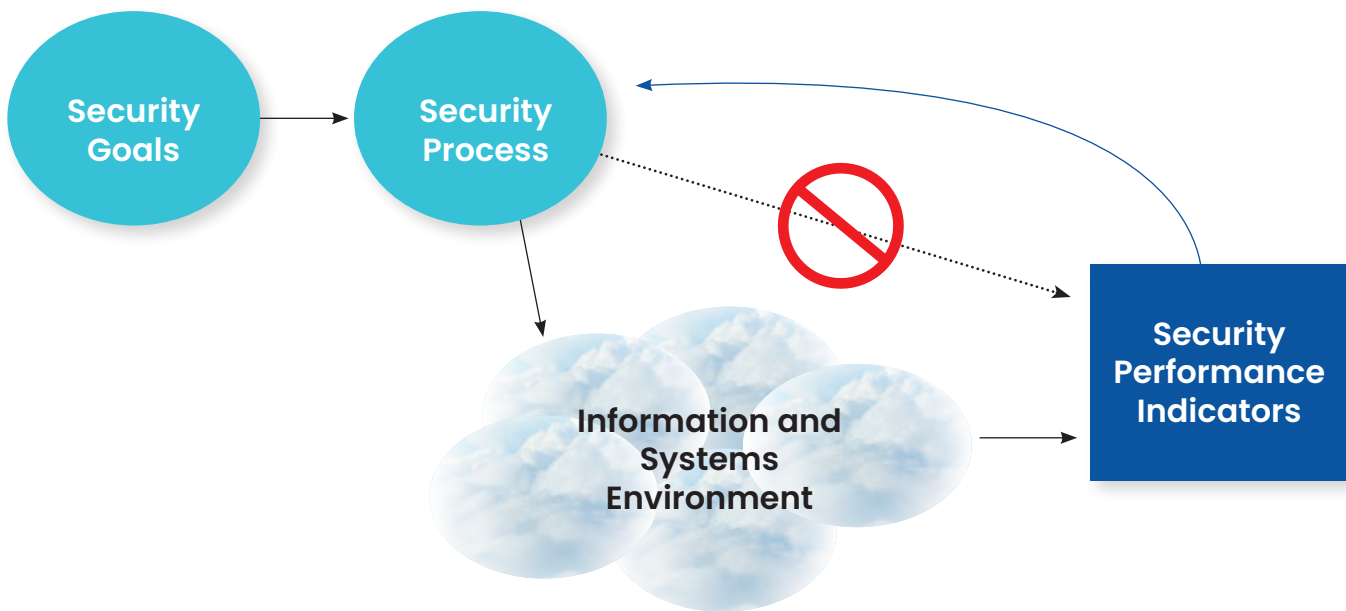


Figure 3: Cybersecurity Measurement Data Flow

Of course, one workflow measurement that misses the target does not imply that such processes should not be measured. One significant source of security measures is an issue-tracking system. Where exceptions to requirements such as security agent installation are detected but cannot immediately be remediated, a process that documents the issue, in combination with the risk and the planned remediation, can be a fruitful source of cybersecurity metrics. Figure 4 shows an issue-register snapshot of identified issues, how severe the risk is if the issue is not addressed, and whether or not remediation plans are executed (and effective). If these snapshots are presented as trends over time, they may provide evidence of both good and bad security program performance.




Figure 4: Issue Metrics

In order to create a trustworthy issue classification system, there must be a connection to actual cybersecurity risk assessment based on business risk appetite. That begins with an authoritative qualitative description of the amount of risk a firm is willing to accept with respect to a given category of negatively impacting events—in this case cyber. Cybersecurity policy, process, standards and procedures must fall in line in support of the risk appetite, and all aspects of a cybersecurity program should be measured for performance.

In this sense, a performance measure may be a *risk indicator*, though still not an indicator that risk is reduced because that requires a demonstration of goal achievement—a validation metric. Goal achievement is measured not with reference to the cybersecurity controls, but rather via independent “sanity checks”—both planned (e.g., breach and attack simulation) and unplanned (e.g., actual breaches).

The extent to which both performance and goal measures accurately reflect the cybersecurity program is a direct reflection of how well it is managed. Note that the information that the metrics provide may show that cybersecurity itself is poor. Even a well-managed program may operate under constraints that prevent it from achieving its goal. But a CISO will not go to jail if all of the CISO’s documentation, including metrics provided to external auditors and investigators, accurately reflects the status of the cybersecurity program’s performance and goal achievement. The internal management debate is then reduced to whether the program is truly delivering risk reduction to a level below management’s risk appetite.

The quantitative version of risk appetite is risk tolerance. Figure 5 is a simplified version of its composition, typically a combination of cybersecurity program goal and performance measures trending over time, collectively called “key risk indicators” or “risk tolerance metrics.” Thresholds should set a theoretical ceiling on where it seems reasonable that risk tolerance trends indicate a breach of qualitative risk appetite. Where the thresholds are breached, postmortems provide an opportunity for systemic practice improvement, including critical evaluation of methods and assumptions.

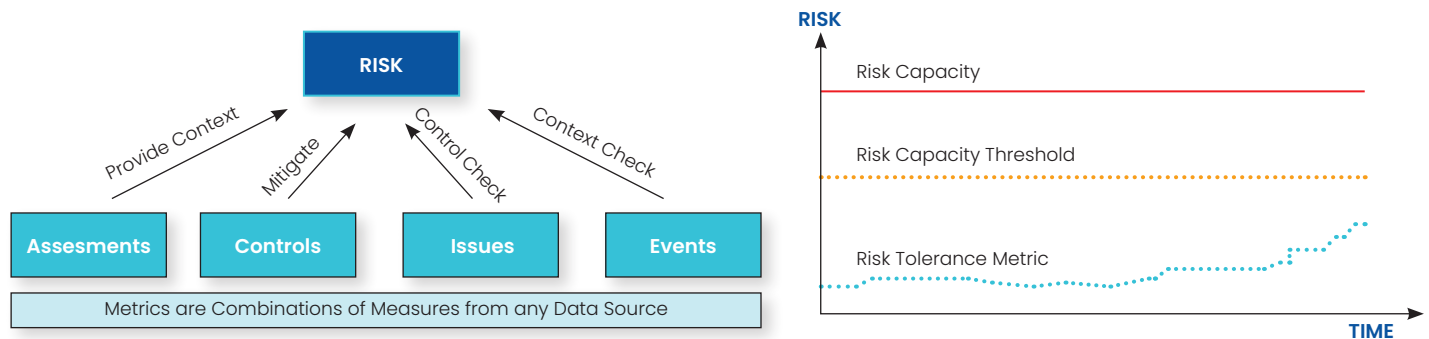


Figure 5: Risk Metrics

In summary, a set of good cybersecurity metrics is an indicator of good cybersecurity management, but neither of those is the same as good cybersecurity. Good cybersecurity metrics often reflect poor cybersecurity despite the best efforts of cybersecurity management. This is a situation similar to other fields where there is an uncontrollable threat (e.g., firefighting, drug counseling, military service). Although there are a plethora of cybersecurity metrics, the key to a good metrics program is completeness with respect to performance metrics, realism with respect to goal metrics, and integrity with respect to both.

This article is adapted from a forthcoming book by Dr. Bayuk: "FrameCyber: How to Reduce Cybersecurity Risk."



USING TIME PATTERNS TO PREDICT FUTURE CYBERCAMPAIGNS

DR. EDWARD AMOROSO

By extrapolating the average time between initial cyber skirmishes and their corresponding full-out attack campaigns, disturbing predictions can be made about future industrial control system attacks, artificial intelligence misuse and global cyberwar.

USING MODELS TO PREDICT ATTACK CAMPAIGNS

During the past quarter century, a pattern has emerged in which some new cyberattack method is demonstrated to work in the wild and, after a period of relative calm, fully exploited at scale roughly 13 years after the initial view. This broad pattern applies to worms, distributed denial of service (DDOS) attacks and attack ransomware.

Using simple extrapolation, it is possible to make predictions about future attack campaigns at scale, based on initial observations currently experiencing relative calm. Specifically, disturbing predictions can be made about industrial control system (ICS) attacks, artificial intelligence (AI) misuse, and global cyberwar.

MODEL 1: WORMS

The first worm¹ was observed in 1988: the so-called Morris Worm. In the ensuing years, worms were certainly known, but it was not until 15 years later, in 2003 that the method was deployed at scale. During that year, the SQL/Slammer, Blaster, Nachi and Sasser worms were unleashed against global infrastructure.



Figure 1. Worm Pattern

MODEL 2: DDOS

The first recognition of the DDOS threat came via warnings from the U.S. federal government in advance of the Y2K transition. Serious DDOS attacks followed in March 2000 targeting CNN, eBay and others. After a relatively quiet period of 12 years, a full unleashing of DDOS fury was aimed at U.S. online banks in 2012, presumably from a nation-state actor.

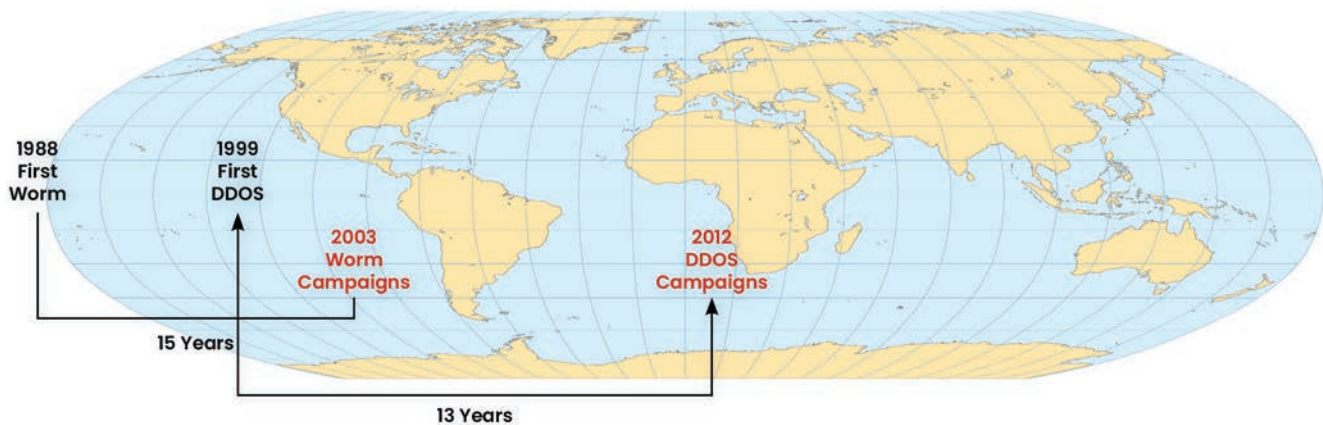


Figure 2. Adding DDOS Pattern

MODEL 3: RANSOMWARE

The first evidence that cryptocurrency could be used for illicit purposes emerged in 2008 with the famous Bitcoin paper. After a period of unease with cryptocurrency, including isolated issues such as Silk Road, the first broad exploitation emerged with ransomware attacks, which reached a peak in 2020 (and continue today).

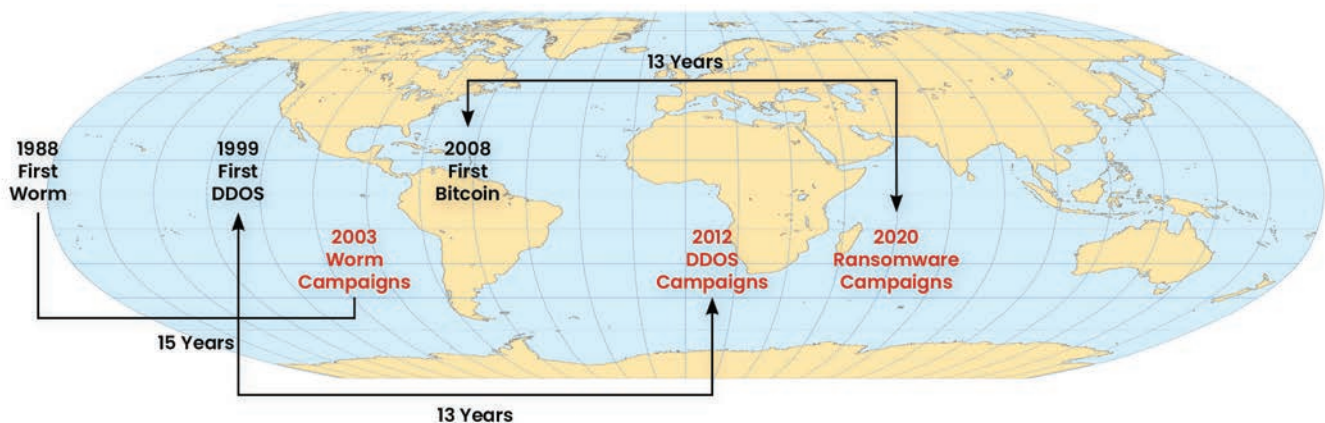


Figure 3. Adding Ransomware Pattern

MODEL 4: ICS ATTACK

The first serious industrial control system (ICS) attack of any real consequence occurred in 2010 with the famous Stuxnet incident, which targeted Iranian nuclear systems. Extrapolating forward, one can predict that a series of disturbing ICS attacks is likely to occur in the coming year, possibly in 2023. Citizens should expect to see hits to factories, power systems, and so on.

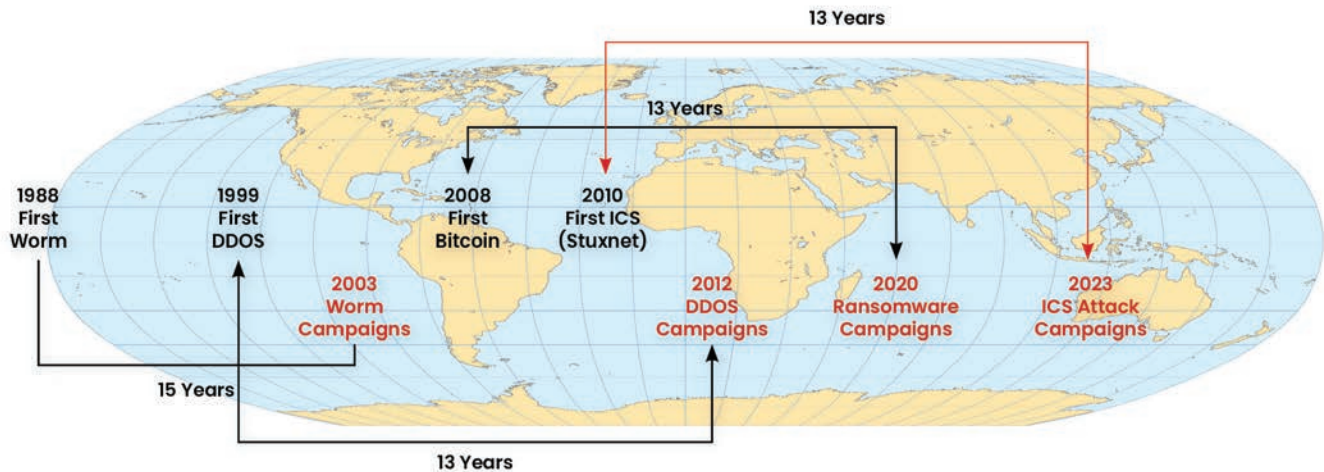


Figure 4. Adding ICS Attack Pattern

MODEL 5: AI MISUSE

The first evidence that AI could be applied to cybersecurity became evident around 2013 with the emergence of companies like Cylance. While this is a benign initial view, one can easily extrapolate misuse of AI to emerge at scale in roughly 2028, which is 15 years after the first occurrence. Citizens should expect to see AI offensive weapons that use AI models for attacks.

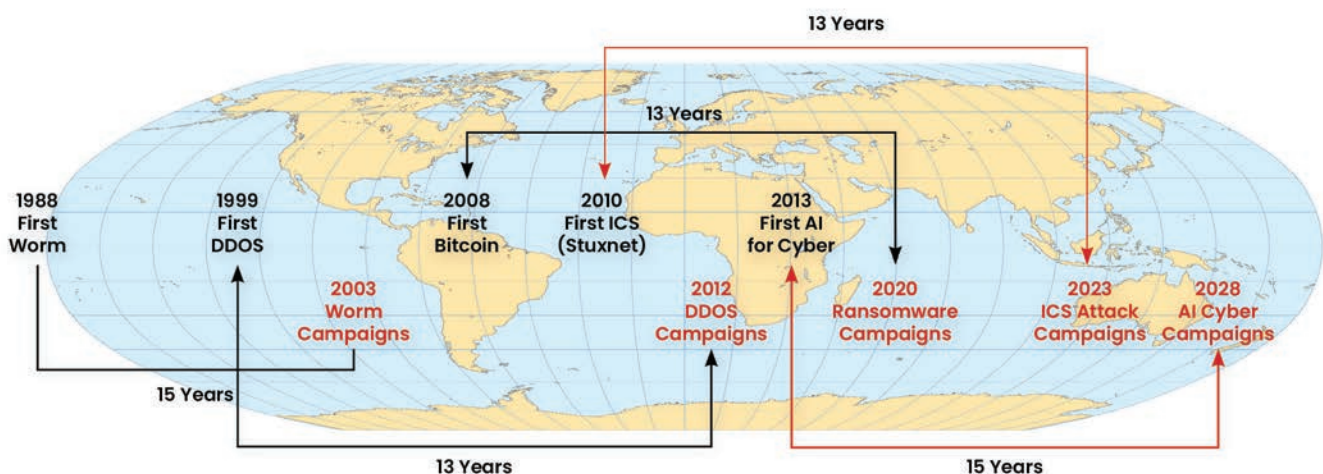


Figure 5. Adding AI Misuse Pattern

MODEL 6: GLOBAL CYBERWAR

The current situation between Russia and Ukraine is more than likely to cascade into a major cyberwarfare situation where the goal is serious cyber dominance, versus making a political or philosophical statement. Extrapolating this geopolitical conflict using our pattern model puts the first global cyberwar 14 years later, in 2036.

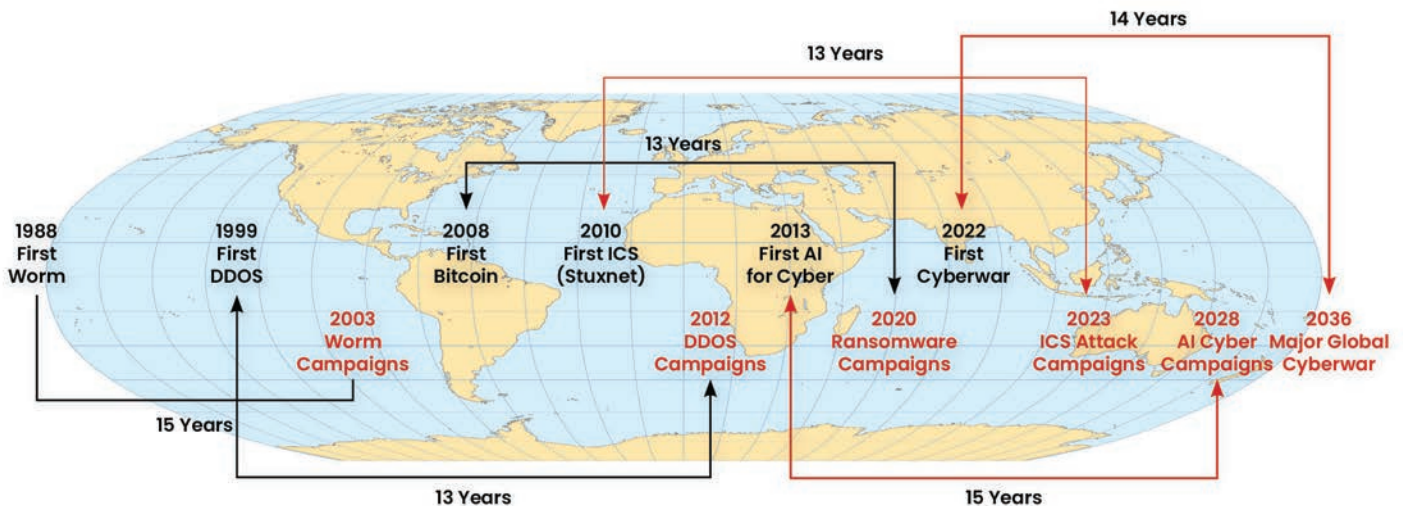


Figure 6. Adding Global Cyberwar Pattern

IMPLICATIONS

Readers will note that no interpretation is made here beyond the simple pattern matching and extrapolation done based on previous and existing data. Nothing about the predictions of ICS attacks, AI misuse and global cyberwarfare should raise an eyebrow for any expert observer. All of these possibilities seem high, and we should view such campaign predictions as grave.

¹ Readers might quibble with the author's designation of what was actually the first observation of a given attack method. Every effort is made to select prominent, meaningful first observations that a given method can work in the wild. Usually, if some other exploitation would have been selected, its emergence date is sufficiently adjacent as to not change the average 13-year thesis proposed here.



Island

Island is the browser designed for the enterprise that makes work fluid, yet fundamentally secure. With the core needs of the enterprise embedded in the browser itself, Island enables organizations to shape how anyone, anywhere works with their information, while delivering the Chromium-based browser experience users expect.

Island, the Enterprise Browser.

TAG CYBER
DISTINGUISHED

REPRINTED FROM THE TAG CYBER SECURITY ANNUAL

©TAG CYBER 2023