TAG Cyber

Security Annual

2023

SPECIAL REPRINT EDITION

HUMAN SECURITY: DISRUPTING DIGITAL FRAUD AND ABUSE WITH MODERN DEFENSE

AN INTERVIEW WITH GAVIN REID, CISO AND HEAD OF THE SATORI THREAT INTELLIGENCE TEAM AT HUMAN SECURITY

THERE'S MORE TO DEEPFAKES THAN MEETS THE EYE

SECURITY METRICS SOMETIMES MISS THE POINT

TAG CYBER DISTINGUISHED VENDOR

HUMAN

The need to reduce cyber risk has never been greater, and HUMAN has demonstrated excellence in this regard. The TAG Cyber analysts have selected HUMAN as a 2023 Distinguished Vendor, and such an award is based on merit. Enterprise teams using HUMAN'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

HUMAN SECURITY: DISRUPTING DIGITAL FRAUD AND ABUSE WITH MODERN DEFENSE
AN INTERVIEW WITH GAVIN REID, CISO AND HEAD OF THE SATORI THREAT INTELLIGENCE TEAM AT HUMAN SECURITY
3

THERE'S MORE TO DEEPFAKES THAN MEETS THE EYE
David Hechler
6

SECURITY METRICS SOMETIMES MISS THE POINT
John J. Masserini
11

REPRINTED FROM THE TAG CYBER SECURITY ANNUAL

TAG CYBER HUMAN

AN INTERVIEW WITH GAVIN REID, CISO AND HEAD OF THE SATORI THREAT INTELLIGENCE TEAM AT HUMAN SECURITY

HUMAN SECURITY: DISRUPTING DIGITAL FRAUD AND ABUSE WITH MODERN DEFENSE

Trying to weed out actual humans from online bot traffic can be a tricky business that has major consequences for security teams and the overall company. Bots can cause serious damage to an enterprise's reputation and bottom line through account theft and payment fraud, as well as fake account creation, reviews and comments.

Tracking over 20 trillion digital interactions each week, HUMAN Security offers a suite of products that prevents digital attacks, bots, fraud and account abuse. To make things easier, it does all the above with just a single line of code. We sat down with HUMAN to get an overview of their products, as well as the advantages they bring to businesses in advertising, marketing, government, education, e-commerce and enterprise security.

TAG Cyber: You started out in the back of a science fiction bookstore, tell us a bit more about your early days and your growth into a market leader.

HUMAN SECURITY: We have been protecting enterprises from digital fraud and abuse for over a decade. Originally based in the back of a sci-fi bookstore, we were founded by Tamer Hassan, Michael Tiffany, Dan Kaminsky and Ash Kalb with the mission to protect the integrity of the internet by disrupting the economics of cybercrime. Over the years, hackers have learned to deploy bots that are so advanced they're practically unstoppable. They're infiltrating companies, taking over accounts, creating fake ones, scraping websites for information and impacting transactions. If that wasn't bad enough, they're also using infected devices and sending fake requests to target websites and apps to steal money and disrupt operations.

Today, HUMAN Security verifies the humanity of trillions of digital interactions each week across billions of devices for more than 450 top enterprises and internet platforms. Thanks to our visibility across the internet, HUMAN is in the position to disrupt digital fraud and abuse through the continuous adaptation of dynamic network, device and behavioral signals. Furthermore, our Satori Threat Intelligence team performs takedowns and disruptions. Examples of our major takedowns of cybercriminal operations include: **3ve, Pareto, Scylla** and, most recently, **VASTFLUX.** All these takedowns have one thing in common: collective protection. Instead of companies and teams individually trying to protect themselves, we protect them all with our Human Defense Platform.

We predict that in 2023, companies will begin to band together to strengthen their defenses and take a stand against digital fraud and abuse.

TAG Cyber: What are the differences between human and non-human cybercrime that businesses need to be aware of when protecting themselves?

HUMAN SECURITY: Behind every cybercrime is a human. Whether they're using a sophisticated bot to execute the crime or not, we're dealing with cybercriminals trying to game enterprises at scale and make as much money as possible with as little cost or risk as possible. Over 77% of digital attacks use sophisticated bots to scale and obfuscate the attack path. For example, cybercriminals benefit from economies of scale by automating the verification of stolen credentials. While non-human and human attack vectors necessitate different detection and countermeasures, businesses can boost their security by fortifying apps, along with landing, login, transaction, checkout, and review pages by ensuring they are engaging with real humans.

TAG Cyber: Describe the key components of your modern defense strategy against bot attacks and fraud.

HUMAN SECURITY: Our technology, processes and relationships have been purposely designed to disrupt the economics of digital fraud and abuse by increasing the cost to cybercriminals, while also reducing the cost of collective protection. We call this "the modern defense strategy." Our visibility in the market is a key differentiator. Today, we verify 20 trillion interactions a week across a total of three billion devices monthly, enabling HUMAN to detect fraud and abuse with unparalleled scale, speed and precision. Our network effect is the feedback loop of technical evidence from up to 2,500 network, device and behavioral signals parsed through 350 algorithms looking for signs of digital fraud and abuse at the time of interaction. Our disruptions and takedowns are led by HUMAN's Satori Threat Intelligence team, which I lead. The team uncovers, reverse engineers and takes down digital fraud and abuse-driven threats. This stops the whack-a-mole process; when we disrupt or takedown a cybercriminal organization, their fraud and abuse go to zero for good.

TAG Cyber: Could you briefly list the various products you offer, as well as their main features?

HUMAN SECURITY: Our Human Defense Platform comprises a suite of products to protect organizations from digital fraud and abuse use cases, while our **Account Defender** product safeguards an organization's app and website accounts by detecting and neutralizing compromised and fake accounts. The HUMAN **Bot Defender** solution protects websites, mobile apps and APIs from automated attacks carried out by sophisticated bots. Next, there is our **Credential Intelligence** product that detects and stops the use of compromised credentials on websites and mobile apps in

real-time. To identify high risk PII, PCI and vulnerability incidents so response teams can act fast, **Code Defender** is a client-side web application security solution that provides comprehensive real-time visibility and granular control into a modern website's client-side supply chain attack surface. To stop marketing campaign fraud, we offer **BotGuard for Growth Marketing** that protects data pools from contamination by preventing sophisticated bots from converting on landing pages. We also offer **cleanAD**, an on-page, behavioral malvertising-prevention solution that protects publishers and platforms from digital attacks executed through the advertising ecosystem. Finally, to protect the programmatic advertising ecosystem and shield it from fraud, we offer **MediaGuard**, thereby improving quality and trust in the digital ad ecosystem.

TAG Cyber: What is the top cyber threat facing companies in 2023?

HUMAN SECURITY: Today's attackers are constantly upping their game to bypass a company's defenses. We're hearing from our clients that account takeovers, fake account creation, and web scraping attacks are becoming more prevalent, as attackers utilize automation to increase their level of sophistication. As a result, it's becoming harder for companies to distinguish between a human and a malicious entity. 'Digital fraud and abuse techniques that easily get past WAFs, CDNs and CAPTHCHAs, so ensuring you have the right protection is critical. That's where companies like HUMAN can help with modern defense and collective protection. We predict that in 2023, companies will begin to band together to strengthen their defenses and take a stand against digital fraud and abuse.

THERE'S MORE TO DEEPFAKES THAN MEETS THE EYE

DAVID HECHLER



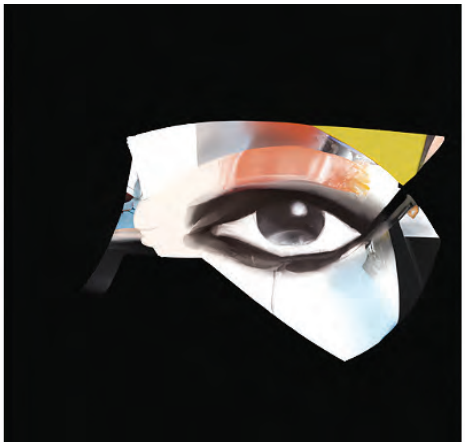

TAG CYBER/DALL-E

What do you think of when you hear the word "deepfakes"? A video featuring Tom Cruise saying and doing silly things? A series of photographs with a face morphing from male to female? A clip of Kim Jong-un in which he addresses the American public? A guy who used to post on Reddit?

Some of you may be hearing (or seeing) that word for the first time. Others know a lot about it. They know that it got its name from a guy who used it on Reddit. And they've seen lots of Tom Cruise memes. They understand that, even though many people think immediately of videos, there are also deepfake audios. And I didn't even mention those, or pornography, in the paragraph above. So you see, there's a wider variety of deepfakes than some people realize.

Let's start with the basics. As the term is understood today, it combines **"deep learning"**—a kind of machine learning—and "fakes." What you're seeing or hearing is not the real thing: Deefakes are built from manipulated sounds and/or images. But the motives behind the manipulation are not all the same. That's why they shouldn't all be lumped together.

THEY'RE NOT ALL BAD

Deepfakes have a bad reputation. The ones that get the most attention are those in which the content manipulators do not ask the people featured in the fakes for permission to use their voices or images, and their motives may be malicious or indifferent to how the individuals affected may feel. But lots of deepfakes are created for amusement and seem harmless. They may be satire or parody. Others are designed to make a serious political point. And many harbor no intent to deceive.

Kim Jong-un deepfake

REPRESENTUS

In fact, some deepfakes announce themselves as fakes. For instance, the **Kim Jong-un clip**, above, was created by the nonpartisan, nonprofit **Represent Us** as a public service ad. The North Korean leader, seated at a desk and clad in a Mao jacket, calmly warns American voters that he doesn't have to work to destroy their country. He points to their partisan divisions and ferocious fights over elections. "It's not hard for democracy to collapse. All you have to do," he says, pausing to crack a smile, "is nothing." The film ends with these words on the screen: "This footage is not real, but the threat is."

Another public service spot used a **deepfake of Joaquin Oliver**, a Stoneman Douglas High School student who was killed in the Parkland, Florida, shooting. His parents introduced him by explaining in a video that he'd been gone for two years and had missed his first opportunity to vote in an election. Now artificial intelligence has allowed him to speak again. The deepfake video of their son follows, and he offers an impassioned plea for people to vote "because nothing's changed, people are still getting killed by guns." He urges them to vote "because I can't."

The many deepfakes of **Tom Cruise** make lighthearted fun of the actor, but in recent years actors have benefited from this new technology. When a documentary about the career of Val Kilmer was being filmed, the actor was not able to sit for an interview because an operation to treat his throat cancer had left his voice badly damaged. But a company called Sonatic has been able to recreate his voice in a way that has **extended** his acting career.

Then there's Bruce Willis, whose health problems led him to retire from acting. But he recently **made a deal** to allow a company called Deepcake (that's not a typo) to map his face onto the body of another actor for a **commercial**. Though there was some disagreement about the circumstances, the message Deepcake was

Joaquin Oliver deepfake

CHANGE THE REF

announcing was clear. As was the company's aim to launch a new industry. Actors who can no longer act, the company seemed to be saying, or actors who have a commitment to perform that conflicts with another opportunity elsewhere, can now digitally clone themselves by authorizing deepfakes.

GRAY AREAS

A director's failure to alert viewers that a voice was deepfaked in a recent documentary stirred controversy.

Some uses of deepfakes have been criticized on ethical grounds for failing to inform the audience. A noteworthy example involved a documentary about Anthony Bourdain that was filmed after he committed suicide. The director had access to thousands of hours of video and audio from his subject's popular food and travel television shows. But in three instances the director wanted to introduce sentences that Bourdain had written but had not recorded. So he decided to use deepfaked audio of Bourdain's voice.

When director Morgan Neville first acknowledged what he'd done, **several critics were aghast**—both that he'd done it and hadn't disclosed it in the film. I can't help but think that it won't be long before people simply accept such things, now that this is an option. I can imagine a far greater uproar had Neville inserted Bourdain deepfaked on video, but this, too, is easy to do. It seems bound to happen. And my guess is that it won't take long before the novelty, and ethical qualms, wear off.

By contrast, there was no need to issue a disclosure when Carrie Fisher and Peter Cushing made deepfaked appearances in "Rogue One: A Star Wars Story." They'd both been gone for years, of course. And one can be sure the use of their images was authorized. Somehow it seemed quite natural, given that this was a science fiction movie, after all. Now the **question** seems to be whether the Star Wars franchise will bring back Fisher, Mark Hamill and Harrison Ford for a deepfaked reunion—deepfaked to make them all youthful again, even though two are still alive. The money seems to say yes, and you can be sure that ethics won't stand in the way.

THE DARK SIDE

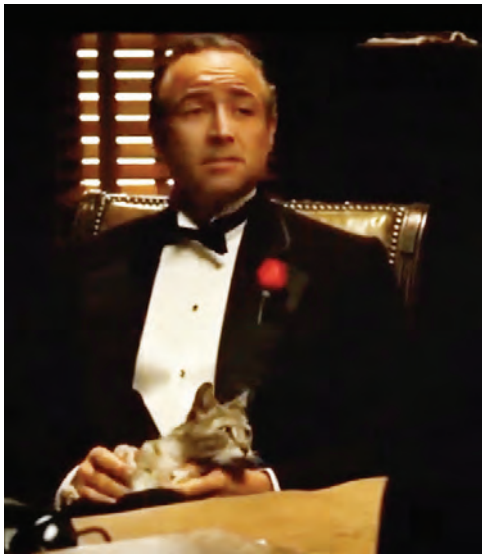

Nicolas Cage as Marlon Brando deepfake

As I noted earlier, the deepfakes that get the most attention are controversial. Obvious examples are the ones created by the Reddit user whose handle gave the concept its name. In late 2017, he began posting on Reddit pornographic videos in which the women's faces had been replaced by those of well-known actresses and other celebrities. As the popularity of his postings grew, he started a so-called Subreddit called deepfakes in which other registered users (known as Redditors) shared their own creations. In addition to pornography, Redditors posted deepfakes of other kinds of entertainment. A particularly popular series which became a genre unto itself offered deepfakes of **Nicolas Cage**. These were often compilations of brief movie clips in which Cage's face was swapped into the bodies of well-known actors and actresses ranging from Marlon Brando in a scene from "The Godfather," to Julie Andrews walking in the hills above Salzburg singing: "The hills are alive with the sound of music." Nothing dark or gray there. Unlike the hard-core content it was paired with, these were just silly.

The Deepfakes Subreddit was eventually shut down, and it wasn't because of the Cage videos. The network **banned** the Subreddit for violating its content policy, "specifically our policy against involuntary pornography," the announcement said. Deepfake pornography is still widely available elsewhere, of course. By at least one measure, it completely dominates the field. In 2019, an Amsterdam-based organization called Deeptrace issued a **report** that found that 96% of all deepfake videos online were pornographic.

Political deepfakes can pose grave dangers if they fool the public, but they're more easily defused because they're out in the open.

To put the Subreddit takedown in context, the unauthorized posting of pornographic images of women by men had been a serious problem since at least 2010. (These earlier postings did not involve deepfakes, but they paved the way for the Deepfakes Subreddit.) It was 2010 when Hunter Moore, from Woodland, California, started isanyoneup.com, the internet's best known **"revenge porn"** website. Moore encouraged people to submit real sexually explicit photographs of women without their consent, which he then posted on the site. They were often supplied by men who bore a grudge. California passed a **law** in 2013 making it crime to post this material knowing that it would cause the women emotional distress, and two years later **Moore** pleaded guilty and was sent to prison. In 2014, the "**Celebgate**" scandal broke in which at least five men hacked into the computers of more than 200 celebrities, including actresses Jennifer Lawrence and Mary Elizabeth Winstead, to steal nude photographs and other private material.

In the years that followed, technology made it easy for anyone to create deepfakes. By 2018, anyone could create them using software programs that were readily available. A short time later, celebrity deepfake videos were easy to create from a mobile phone.

PLAYING FOR HIGHER STAKES

Some of the most dangerous deepfakes have been ones that have targeted political leaders. The danger was in the potential consequences if they had been believed. During the U.S. presidential campaign in 2020, some videos promoted by the Trump campaign appeared to show Joe Biden as old, tired, confused and out of touch, but they **were actually deepfakes.**

Nearly two years later, Russia was engaged in a different kind of campaign. Three weeks after the country invaded Ukraine, a **deepfake of Ukraine President Volodymyr Zelensky** was broadcast showing him addressing his soldiers and instructing them to lay down their arms. The video was promoted by Russian social media along with posts on Facebook, Twitter and YouTube. In both instances, the targets quickly called out the fakes and they were removed from wide distribution. In Ukraine, the government had even warned its citizens in advance to expect Russia to engage in this kind of subterfuge.

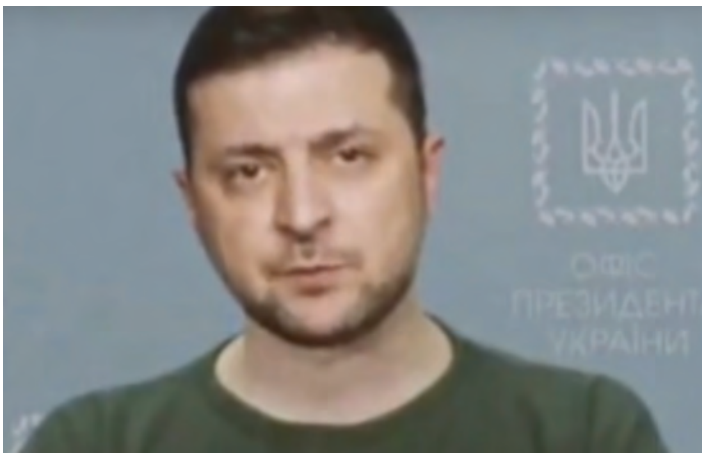

Volodymyr Zelensky deepfake

As serious as those incidents were, in one important respect they were easier to defuse than many other deepfakes for one simple reason: They were out in the open. That was the whole point. They were designed to influence public opinion. But that

also meant that they were closely scrutinized by journalists and experts of all stripes. It didn't take long to identify what they really were.

By contrast, criminals thrive on stealth. They often use deepfakes to try to trick businesses into wiring them funds, or they extort money by threatening to expose the image of a CEO in a compromising position. And companies are often reluctant to reveal anything about these episodes—whether they succeeded or failed, whether the images were genuine or phony—for fear of tarnishing their reputations. So it can be hard to know how big a threat deepfakes represent.

One indication that it's growing can be found in VMware's annual Global Incident Response Threat Report. In June 2022, it surveyed 125 cybersecurity and incident response professionals and found a 13% uptick in deepfakes year over year. And 66% of respondents had seen them during the previous 12 months, with email cited by 78% as the most common delivery method.

HELP NOT WANTED

This technology is new enough that innovations seem to pop up regularly. Here's a new twist. Now that so much work is conducted from remote locations far from traditional offices, it's no longer unusual for job interviews to be conducted remotely, and for employees to work for years for bosses they haven't met and may never meet. So perhaps it shouldn't be shocking that some companies have found they've hired not the fine young man or woman they thought they had, but a deepfake instead.

Last June, the FBI issued an **alert** that warned companies about deepfake job candidates. Complaints along these lines have been growing, the bureau noted. Rick McElroy, principal cybersecurity strategist at VMware, said it shouldn't be surprising. As companies have improved their security, criminals looked for other ways to break in. "Organizations have spent an inordinate amount of money on these controls," he said. "Manipulation of the human is the easiest way—it's the fast forward button."

Humans have even supplied the raw materials the criminals use to create deepfakes. We give them up ourselves when we post photos, videos and audio files on websites and on social media. And the ability of technology to turn stolen identities into deepfakes is improving rapidly. It isn't flawless, McElroy said. The FBI alert noted that audio and video are sometimes imperfectly synched, and that can help companies detect deepfakes. But in the hands of skillful criminals, it's often good enough.

For the criminals, there are real advantages in using this approach, McElroy continued. Human imposters might succeed in securing the same jobs, but they would be hard-pressed to apply for positions at companies around the country or around the world. Deepfakes can scale. And once they obtain employment, they can look for opportunities to steal money if their handlers are criminals, or engage in espionage if their owners are nation-states. (Or do both.)

What strikes me as particularly unsettling is that if you hire and eventually uncover the true "identities" of **deepfake employees**, you may still be left wondering who created them and who they really worked for.

Now that we've explored the wide range of deepfakes—from light entertainment to those that may be most important to consider, but also most unpleasant—this might be a good time to click on one of those "Tom Cruise" videos that you'll have no trouble locating on the 'net. I find they have a welcome calming effect.

SECURITY METRICS SOMETIMES MISS THE POINT

JOHN J. MASSERINI

Before we begin, I'm going to ask for your indulgence for a moment while I share something a bit personal. I know it may seem odd at first, but I promise it will all come together quickly, as will its tie-in with security metrics.

If you've ever met me in person, you would know that I'm a "Big Guy." I'm 6'1" and I go about 240. Now, if we've never had the pleasure of meeting in person, you likely have an image of a fairly round and portly guy, and frankly I don't blame you. My Body Mass Index (BMI) is about 31%, and by every medical definition ever published, I am somewhere between obese and morbidly obese.

The idea behind BMI is that a "healthy" person of a given height should be within a range of weights. It's a well-intentioned effort to give the general population an understanding of what their "optimal" weight should be. But when we look at it closely, BMI is nothing more than a metric used by the medical profession to put some type of measurement on a person's weight/height ratio. Unfortunately, the BMI calculation doesn't consider the type of weight a person carries—whether it's fat, muscle, or water—only that they have it. Because of the lack of context behind the BMI, it can be misleading as a person's true health status. For example, every world-class bodybuilder, who averages 3%-5% body fat, is morbidly obese according to the BMI. Kind of strange, huh?

Why is this important? Well, over the past several years, I have worked incredibly hard to shed a lot of the unhealthy weight I carried. But in doing so, I've packed on a bit of muscle. Since muscle is far more dense than fat, only a little muscle weighs the same as a lot of fat, so looking at my BMI, you wouldn't know that I've dropped almost three pant sizes. And while I can't quite fit in a large, my extra-large shirts have plenty of room now. I am arguably in the best shape I've been in for decades, yet my BMI hasn't changed throughout this journey.

There are metrics that I need in order to manage risk across my enterprise, and there are metrics that my executives are interested in.

Now, I'm sharing all of this to prove an important point that every security executive needs to come to terms with: Even though they are well intentioned, just like the BMI, security metrics can be horribly misleading.

Don't get me wrong. I am a huge advocate of measuring your security program and leveraging those metrics to communicate risk with all of your stakeholders. That said, all too often those metrics are used for shock and awe rather than communicating important messages around risk. I have lost count of the number of meetings I've been in over the years that talk about how many thousands or millions of spam messages were blocked or how many open vulnerabilities there are, but never once mentioned the single phish that got in which caused a department's worth of people headaches for more than a few days. After all, how many times have we seen the fancy PowerPoint deck talking about firewall blocks or packets analyzed, but never anything that speaks to the reduction of risk in the environment.

After countless years as a CISO presenting to boards, executives and colleagues, I've found that I've developed almost a split personality when I'm asked about what metrics to track. There are metrics that I need in order to manage risk across my enterprise, and there are metrics that my executives are interested in. Sometimes they are the same, but most times they are not.

OPERATIONAL VS. RISK METRICS

Whether we like to admit it or not, many of us run the operational side of security as well as the policy or strategic side. When running an operation whose sole focus is defending against attacks, the kinds of metrics I want collected are of very little interest to my board. Do I care about the number of packets analyzed or the number of spam messages blocked? Of course I do. But it's far more about ensuring I have enough headroom with my solution than the amount of risk I mitigate. And more to the point, I am not about to scare my board with fear-inducing, over-inflated numbers that serve no purpose.

Here's an analogy I use a lot. The National Traffic Safety Board doesn't report on how many miles Teslas drive every year, but they certainly report on how many of their vehicles catch fire. The same logic applies to metrics. We don't need to report when our solutions are doing what they are supposed to—only when they don't.

If you feel compelled to talk about the sheer volume and quantity of the statistics you're collecting, do yourself (and your board) a favor and talk about efficacy, not volume. Telling your board that your anti-spam solution is 99.9735% effective means far more to them than saying you blocked a gazillion spam emails. And as a side benefit, you get to open up a dialogue that tells them something they need to hear: No solution is 100% perfect. There you go: a win-win.

When we get down to it, the board doesn't really care about how you run your SecOps. You're the expert they hired, so they expect you to manage what you do. That said, communicating risk to the board is also a critical function of your job, and they expect you to be able to do that effectively. Understanding how your board thinks is critical to your success, but even more important is understanding that they are not security geeks, so developing your metrics program around technical risks is not the best approach.

Your goal is not to use metrics to scare your executives, but to find metrics that they can relate to. To quote one of the most influential psychiatrists of the 20th century, Milton Erickson once said:

"Every person's map of the world is as unique as their thumbprint. There are no two people alike. No two people who understand the same sentence the same way.... So in dealing with people, you try not to fit them to your concept of what they should be."

Ponder that for a moment. Most of us deal with boards and management teams that comprise scores of participants. Your metrics need to make sense not to the one person you are speaking to, but the

dozen or more board members who come from diverse backgrounds and experiences. You don't have one different map of the world to deal with, but dozens—dozens of people who all heard the exact same words you spoke, and who all interpreted those words slightly differently. Well-planned metrics bridge the communications gap that comes with having multiple world maps in your boardrooms.

So, after all that, what are some of the metrics I rely on most? Well, I'm glad you asked. But rather than share specific metrics I like, I think it's more useful to share themes I've found to be highly successful.

Rather than share specific metrics I like, I think it's more useful to share themes I've found to be highly successful.

OPERATIONAL METRICS

Even after all of this, I admit I do share certain operational metrics with my executives and board.

- **SOC Efficacy:** Metrics like Mean Time to Close (MTTC)/Mean Time to Resolve (MTTR) reflect the efficiency of the SOC team in resolving events and closing incidents. This is a key indicator of staffing challenges in the SOC and highlights the potential need for hiring or training existing staff. There are numerous other SOC-related measurements you can identify, so pick the ones that not only measure risk reduction, but also demonstrate value and effectiveness.
- **Compound Annual Growth Rate (of events and incidents):** In the financial world, CAGR is a common term with a well-defined meaning. By using this metric to represent the growth of events, incidents and attacks, the executives understand the reasoning that triggers the budgetary investments required in the security infrastructure and SOC. Used hand in hand with the MTTC metric.
- **Solution Efficacy:** The overall effectiveness of the existing solutions. This is where we measure spam, NIDS/NIPS, antivirus and any other solution we have deployed. This is also used to show the adoption rate of new measures like multifactor authentication, privileged access management and user certification hygiene.
- **Solution Life Expectancy:** This metric shows any security solutions that have less than 20% headroom or are beginning to show a decreased efficiency due to changes in infrastructure, attack vectors or business functions. Primarily used to set the stage for budgets or capital expenses.

RISK METRICS

Ultimately, this is the bread and butter of any metrics program. Each of the categories below can leverage the same data collection for mitigating risks as well as communicating those risks to executives.

- **Attack Metrics:** Attack metrics are arguably the easiest to obtain, the hardest to use effectively and the most susceptible to succumb to the pitfall of shock and awe. Here's the thing about attack metrics: While the month-over-month volumetrics are important, most of the rest of it is useless noise. Are we really at a point where we need to highlight the same port scanner that hits you every month? No, we're better than that. We will talk about the new attack(s) we're seeing that we are susceptible to, and what we're doing about them, but let's not waste everyone's time talking about the attacks that are dropped on the floor because our firewall/IPS is doing its job.

- **Vulnerability Metrics:** The stalwart of the metrics world is undoubtedly reporting vulnerabilities. The key to effective vulnerability metric reporting is to relate them to the potential financial impact on the company. Do not report to the board a count of generic five-tier risks (none through critical) without offering insight into the financial impact of your critical systems. Again, avoid using these numbers to instill fear, but rather, put these findings into context by associating them with the revenue that could be impacted by attacks

- **Identity Metrics:** As more enterprises begin planning their long-term, zero trust initiatives, having a clear understanding of your access controls is critical. Understanding how identities and accounts are created, maintained and ultimately deleted is a foundational necessity when you consider zero trust. Tracking topics such as role ratio, mean time to close, recertification requirements and "out of compliance" metrics will drive a deeper understanding of identity-related risk throughout the enterprise. Also, do not forget to collect and evaluate identity metrics around your AWS/GCP/MSA cloud environments, as access control risks are substantially more risky when you consider most DevOps processes.

- **Availability Metrics:** It seems all too often the availability of a system is prioritized well behind the confidentiality or integrity of a system, rather than giving it an equal footing. Have you done a business impact analysis on that 30-year-old system that runs that old Cobol-68 program which just happens to drive 75% of your revenue? Well guess what? The board wants to know you're on it and there's a plan to ensure it's upgraded, migrated or backed up even though there isn't a published exploit anywhere in the world. If you've forgotten what **C.I.A.** (confidentiality, integrity and availability) is perhaps it's time for a refresher.

- **Regulatory Metrics:** We all have them—whether it's PCI, HIPAA, SOX or any other government/industry related acronym—and regulatory requirements are something we all have to deal with. When discussing these risks with your board, do not just talk about the gaps you have. Make sure you also articulate the potential fines—especially in this GDPR world—and how those gaps could directly impact the levels of fines faced. Again, it's easy to fall into the trap of instilling fear with this, but try to avoid it. Use as much realistic data as possible, especially when dealing with publicly disclosed fines.

So, is your next board meeting going to be filled with fear-inducing, shock-and-awe, BMI-type metrics, or are you going to focus on communicating those risks that the board needs to hear in a way that they can relate to?

Remember, every person in that room interprets your words in their context—not yours. Make sure that your metrics bridge the maps of all the worlds before you.

HUMAN

HUMAN is a cybersecurity company that protects 450+ enterprises by disrupting bots, fraud and account abuse with modern defense. We verify the humanity of more than 20 trillion digital interactions per week, protecting against account takeover attacks, fake account creation, payment fraud, content manipulation, content scraping, PII harvesting and denial of inventory/stockout attacks.

TAG CYBER DISTINGUISHED

REPRINTED FROM THE TAG CYBER SECURITY ANNUAL