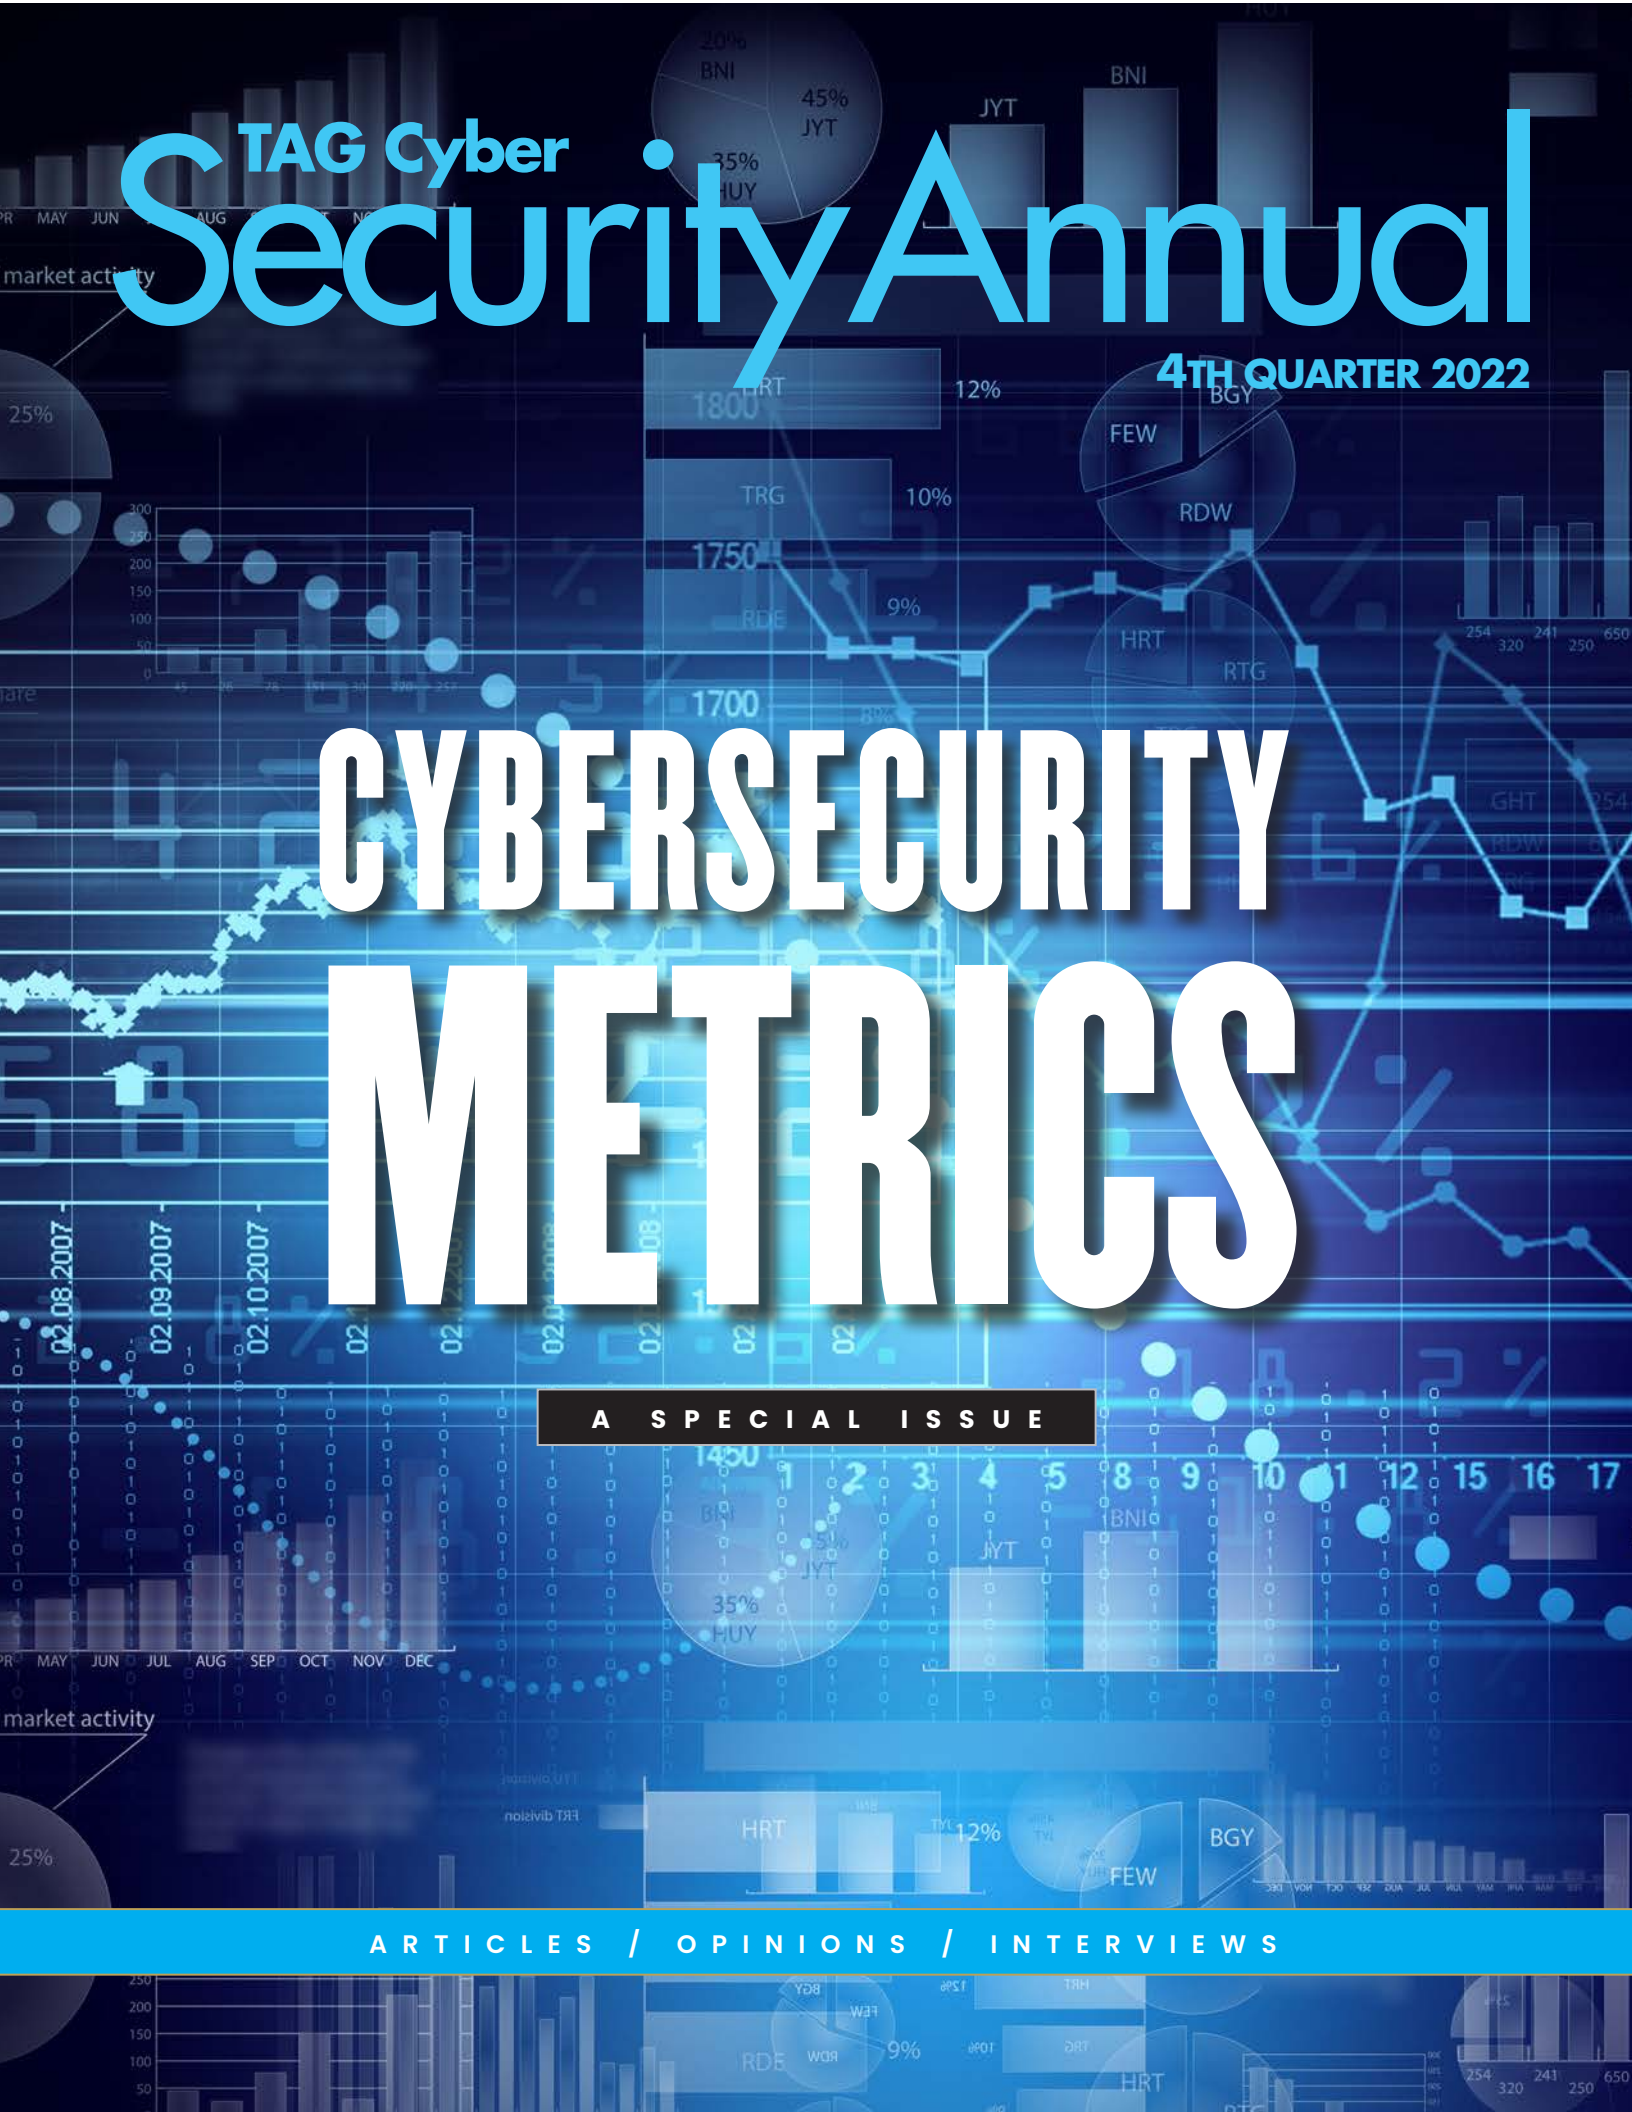

TAG Cyber

Security Annual

4TH QUARTER 2022

CYBERSECURITY METRICS

A SPECIAL ISSUE

ARTICLES / OPINIONS / INTERVIEWS

WELCOME TO THE 2022 TAG CYBER SECURITY ANNUAL 4TH QUARTER EDITION

LESTER GOODMAN,
DIRECTOR OF CONTENT,
TAG CYBER

Dr. Jennifer Bayuk has a passion for metrics. Our resident expert on security metrics, she has a lot to say in a Quarterly chock-full of feature articles devoted to the topic. Dr. Bayuk contributes three very different pieces on the subject herself. Her first tells us what good cybersecurity metrics look like. And if you think you know that already, she'll force you to think again.

She also invited guest contributor Wade Baker to join us., and he makes an immediate impression. After telling us what a "metrics nerd" he is, Baker confesses: "I HATE traditional vulnerability management metrics....Maybe I'm embarrassed about those pretty, yet ultimately pointless, spreadsheets I used to make to 'measure' our risk exposure from unpatched vulnerabilities." He goes on to talk about "metrics that actually matter." And one way or another, so do all of our authors. They all distinguish the useful ones from the useless—or even counterproductive.

Analyst John J. Masserini tells us that metrics used the wrong way can do more harm than good. He cites the example of security metrics deployed to bowl over the board when what directors really need are calm messages about risk that they can understand.

Our leader, Dr. Edward Amoroso, writes a provocative piece on the proposed Bureau of Cyber Statistics, which is provoking a lot of excitement in Washington. If the legislation to create it comes through, Dr. Amoroso says, it may be useful to researchers like us, but the new agency isn't going make us any more safe.

That's just a sample of our metrics that matter. When you scroll down, you'll find much more.

We thank our Research as a Service (RaaS) customers in enterprise, along with our Content as a Service (CaaS) customers in the security vendor community, for providing the support to enable our research and writing. It is through their kind support that we can offer this publication to readers for free.

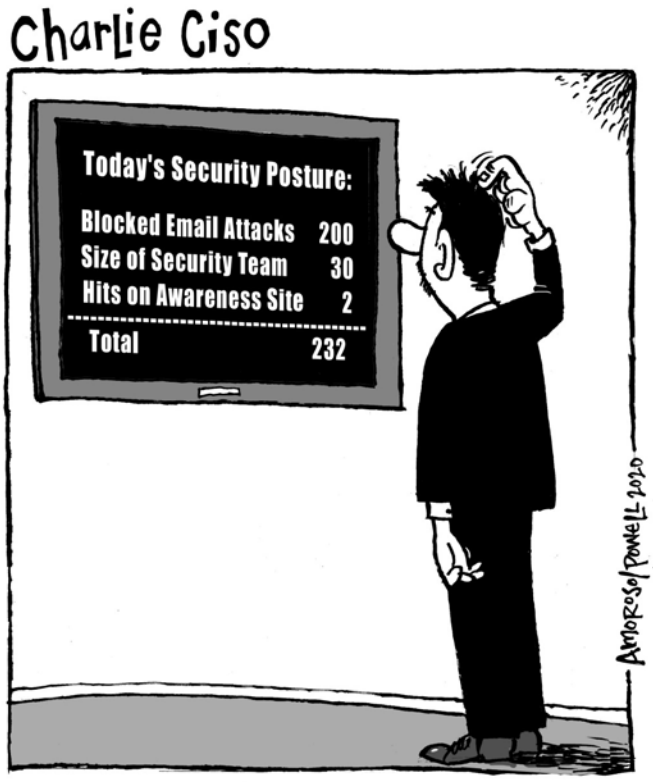

TAG Cyber Security Annual
4TH QUARTER 2022

Lester Goodman, Director of Content

David Hechler, Editor

Contibutors

Dr, Edward Amoroso
Wade Baker
Dr. Jennifer Bayuk
David Hechler
John J. Masserini
Dave Neuman
Christopher R. Wilder

Editorial & Creative

Lester Goodman
David Hechler
Michelle Perino
Managing Editor
Julius Williams
Miles McDonald
Rich Powell

Research & Development

Matt Amoroso
Shawn Hopkins

Sales & Customer Relations

Rick Friedel
Trish Vatis
Laurie Mushinsky

Administration

Liam Baglivo
Julia Almazova

Dr. Edward Amoroso, Founder & CEO

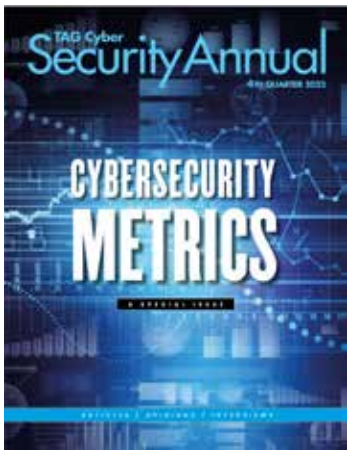

Volume 8 No. 4

October 15, 2022

How Charlie Ciso Cartoons are made!

1: Ed and Rich carefully review 10,000 cartoon ideas from readers around the world.

2: AI tools predict the 300 ideas of the 10,000 that will be best.

3: Ed then develops 300 pages of cartoon copy in handwritten calligraphy.

4: Rich then hand sketches 300 sample pieces of art in colorized panels.

5: Two thousand cyber experts vote to select the BEST cartoon of the 300.

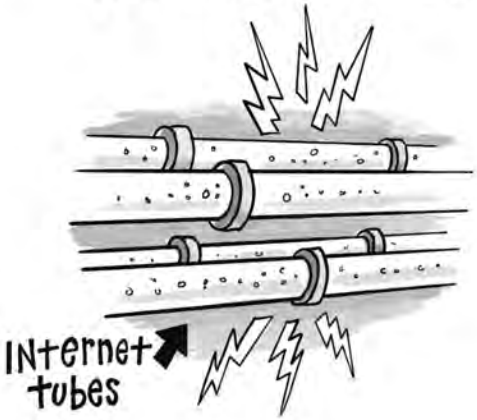

6: We deliver you the final selected weekly cartoon.

Powell/Amoroso

*Meet Charlie Ciso and find out more **here***

FOCUS:

CYBERSECURITY METRICS

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).

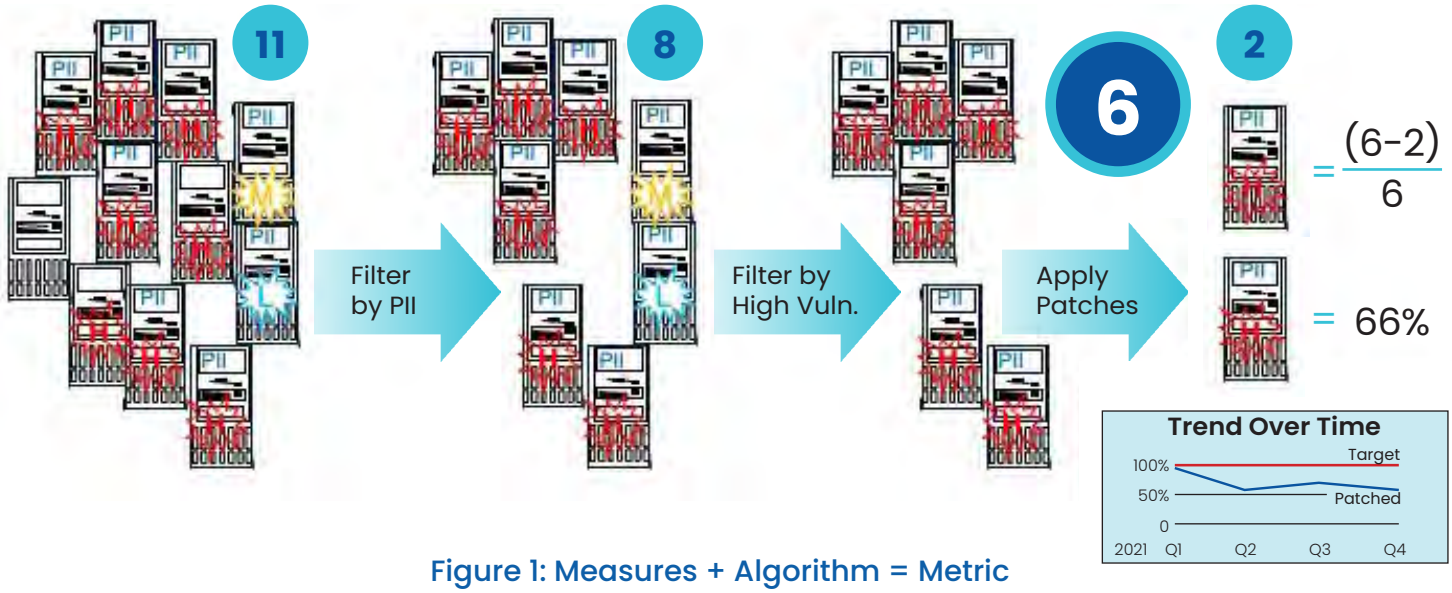

Figure 1: Measures + Algorithm = Metric

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).

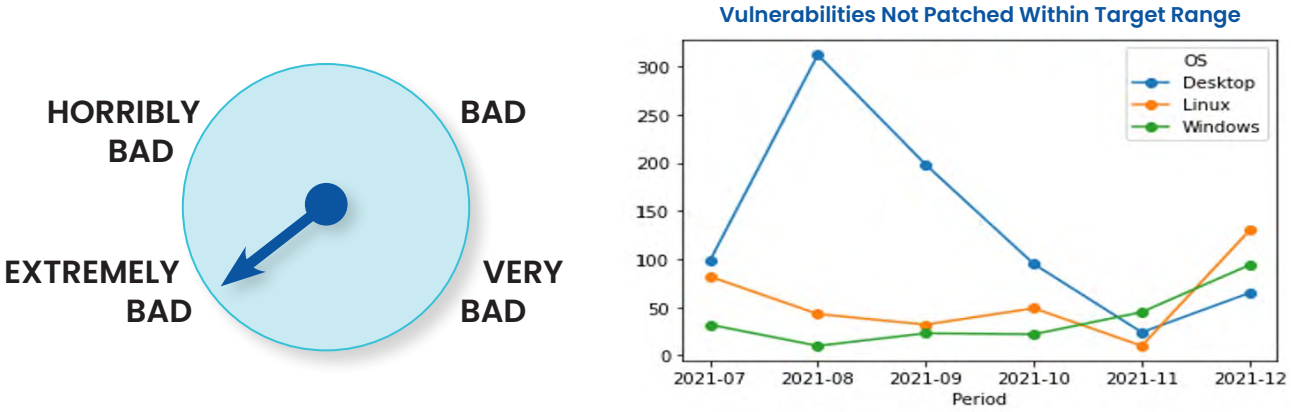

Figure 2: McGraw's Badness-Ometer and an Exemplar Metric

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.

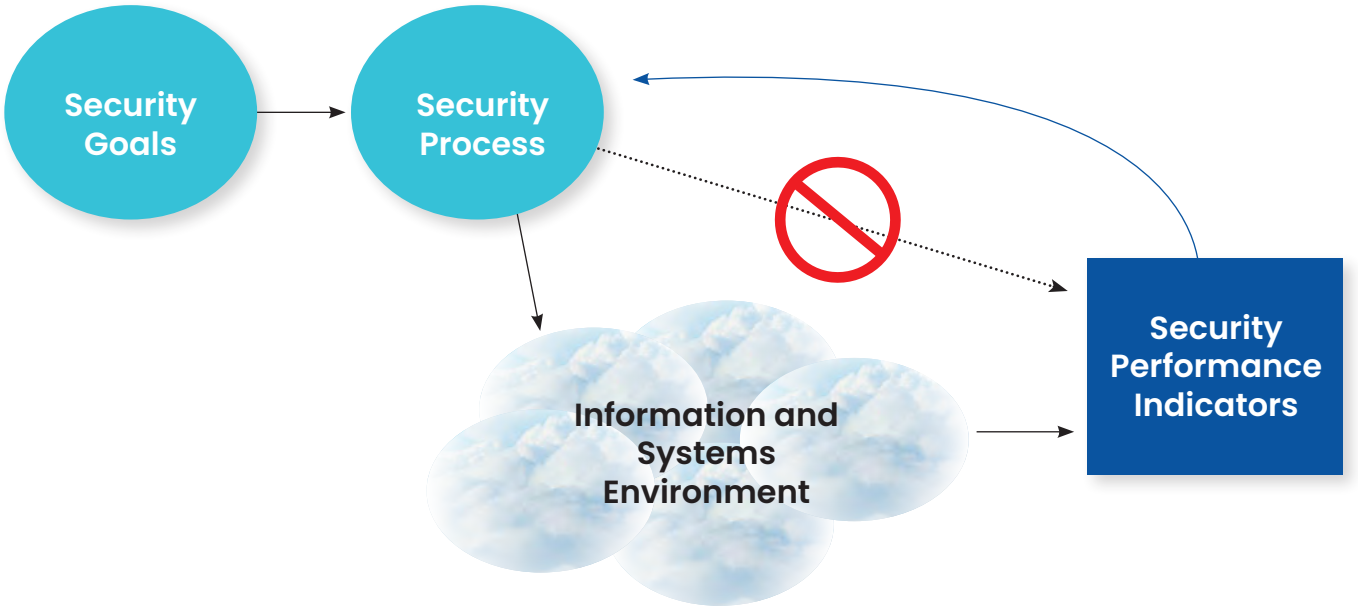

Figure 3: Cybersecurity Measurement Data Flow

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.

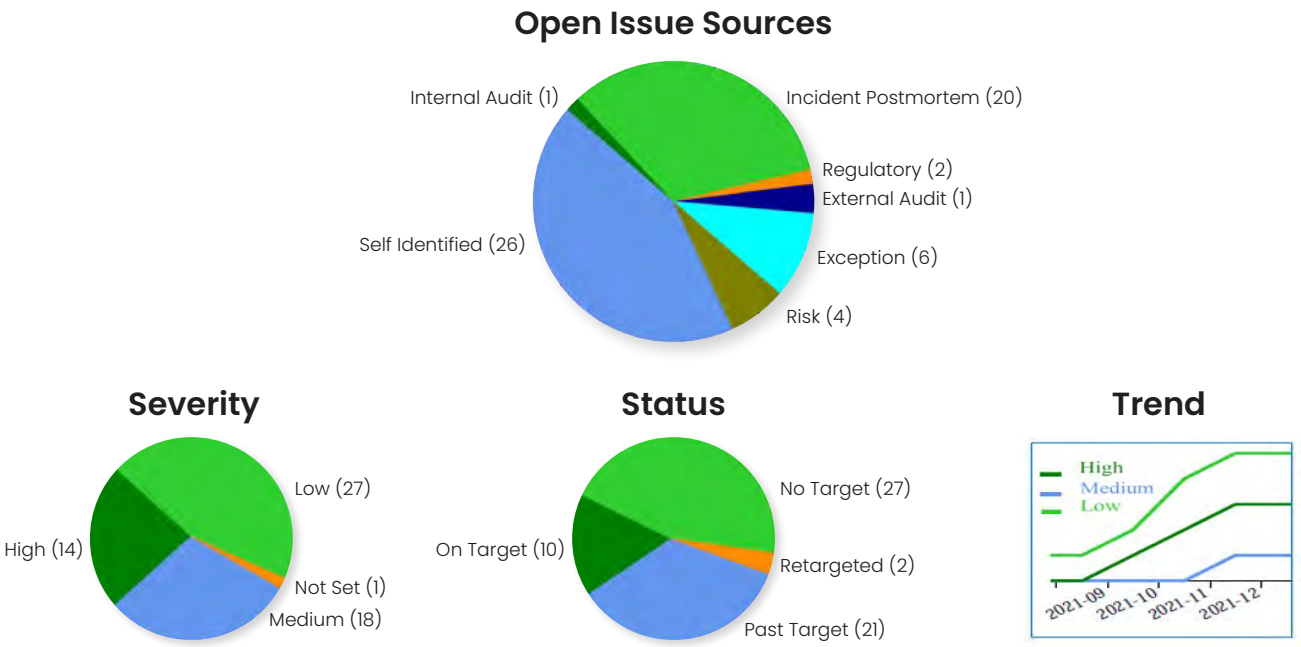

Figure 4: Issue Metrics

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.

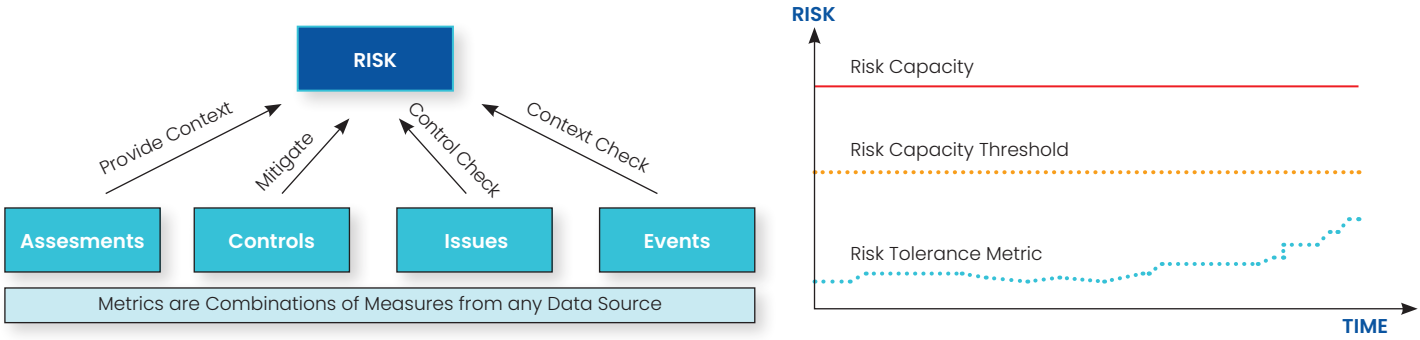

Figure 5: Risk Metrics

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."

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

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

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.

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.

THE PROPOSED BUREAU OF CYBER STATISTICS WILL NOT REDUCE RISK

DR. EDWARD AMOROSO

U.S. National Cybersecurity Director Chris Inglis and many others have actively promoted the idea of a Bureau of Cyber Statistics (BCS). Such a new federal agency would collect and analyze reported data on cyber incidents and crimes, presumably to assist government, industry and researchers.

The **Cyberspace Solarium Commission** is generally referenced as having originally proposed the idea, and this group is filled with experts who know what they are doing. They have released one excellent paper after another—all required reading, in my opinion. And it's hard to dispute anything they would recommend. These are experts.

But while having good cybersecurity data is always desirable, especially to support the type of industry research we do at TAG Cyber, I worry that focusing on a new BCS just reinforces and restates the obvious. Specifically, I'm concerned that the spotlight on a BCS would push our collective attention away from the real security problem: the complexity of our systems.

Here's what I mean. Let's examine the pre and post conditions of having an effective BCS in place. Today, without such a bureau, we are all pretty certain that the effectiveness of managing cyber risk is low. Stated simply: We all know that we're up the creek with no paddle, so to speak. No one disputes this or believes our cyber problems to be minor.

Evidence for such belief is easy to find: just Google it. You'll find Brian Krebs reporting on one incident after another. You'll find the team at Verizon publishing good stats and points on the topic. And you'll find excellent stats and reporting from the FBI. This is depressing reading, but it is consistent and abundant.

I'm concerned that the spotlight on a BCS would push our collective attention away from the real security problem...

Now, if we had a Bureau of Cyber Statistics, a more centralized and presumably accurate repository would exist of reported incidents. And yes, my guess is that it would be even more depressing than the materials we have today. It would certainly reinforce the fact that we have a massive problem. And it would dominate the news cycle from time to time.

Incident reporting entities such as corporations might not like it much—and I tend to agree with them. There is dogma in the federal government that companies actually do know how to stop cyber threats. They just choose not to, and they suppress information related to incidents to avoid embarrassment. This is nonsense, of course. CISOs work like dogs to reduce risk.

So, the obvious question is this: What's the ultimate purpose of the BCS? If the goal is to support researchers (like me), then I am all for the initiative, although this dampens the punch any press release might have. "New agency created to help cyber researchers" will not be breaking news on CNN.

But if the view is that a new Bureau of Cyber Stats will somehow have a meaningful effect on cyber incidents, cybercrime and cyber risk—well, then I do not agree. Collecting more data, albeit more accurate and complete data, would have little or no impact on much of anything, other than our collective depression at our poor security success.

The situation is somewhat akin to a neighborhood experiencing a spike in crime. Everyone in the neighborhood would know darn well what is going on and would demand action. Put yourself in the community meeting where the local assembly person suggests collecting more data on the problem as the solution. Such a proposal would not go well.

When we understand our systems, we can secure them. When we do not understand our systems, then someone smarter can find a way inside.

As I've alluded to above, our cyber problems stem from complexity. For me this is a "period, end of sentence" claim. For over 40 years, I've watched as nearly everything we do in computing and networking has doubled, tripled and quadrupled in complexity. It's been the counterbalance to all the security tools we've bought and deployed.

I thus believe that perhaps the Cyberspace Solarium Commission should have proposed creating a National Bureau of Cyber Complexity Reduction (NBCCR). This could still be metrics-based, perhaps offering new ways to measure complexity (and no—counting lines of code would not be the best approach).

The only hope we have to reducing cyber risk is to reduce complexity. When we understand our systems, we can secure them. When we do not understand our systems, then someone smarter can find a way inside. And the bad guys have an easier job: They only need to find one way in, whereas we need to close all paths of entry.

So, I guess my ultimate view is that the BCS would be a welcome resource for researchers—and our team at TAG Cyber would immediately consume the data. But if the mainstream media, politicians and decision-makers believe that such a bureau would have any material impact on cyber risk, then they are mistaken. And it's our duty to make sure that they understand this fact.

METRICS THAT ACTUALLY MATTER FOR VULNERABILITY MANAGEMENT

WADE BAKER

I'm a self-admitted metrics nerd. I've been collecting data to measure various aspects of cybersecurity management for the better part of two decades now. I love putting numbers on things and monitoring how they change over time. Yeah—I know—get a hobby, right?

But as much I love metrics, I HATE traditional vulnerability management (VM) metrics. Maybe I'm scarred from experiences early in my career when I decried to upper management the number of "critical" vulnerabilities without eliciting a single raised eyebrow. Maybe I'm embarrassed about those pretty, yet ultimately pointless, spreadsheets I used to make to "measure" our risk exposure from unpatched vulnerabilities. It could also be all that time I wasted frantically researching the latest bug promising to bring down the internet (on Friday afternoon, of course) just so we could take our sweet time fixing it.

Anyway… I have issues with how most VM programs go about counting and squashing bugs. Thankfully, I now know there's a better way. Over the last several years, I've had the opportunity to study the vulnerability management practices of hundreds of organizations in great detail through my involvement in a series of research reports titled **"Prioritization to Prediction"** from the **Cyentia Institute** (which advances cybersecurity knowledge and practice through data-driven research, and where I'm a founding partner) and **Kenna Security** (a pioneer in risk-based vulnerability management and now a part of Cisco). Through that research, we've identified numerous ways to measure and improve VM performance. I'd like to share my top three VM metrics with you in this article.

We've identified numerous ways to measure and improve vulnerability management performance. I'd like to share my top three VM metrics with you.

REMEDIATION COVERAGE

Let's start by scrapping that tired old metric for the number of patched (or unpatched) vulnerabilities with a **Common Vulnerability Scoring System** (CVSS) score of X or above. Most organizations track that metric thinking it measures something about their risk exposure, but it doesn't. In fact, if the point is to answer the question "Are we remediating the riskiest vulnerabilities?" that's quite possibly no better than randomly selecting bugs to patch. I know that sounds like hyperbole, so I'll share data to back it up.

First, though, I want to define a metric that I'd place on the top of the shortlist if I were running a VM program: remediation coverage. Coverage measures the completeness of remediation efforts; in other words, what percentage of exploited vulnerabilities were actually remediated.

With that defined, let's look at Figure 1, which compares CVSS-based remediation approaches for vulnerabilities published to the Common Vulnerability Enumeration **(CVE) List**. There are a lot of numbers in the table, so let's focus on the "Coverage" column.

Remediate Above CVSS Base Score	Remediated Correctly (True Pos.)	Remediated Incorrectly (False Pos.)	Remediated Too Soon (False Pos.)	Delayed Correctly (True Neg.)	Efficiency (Precision)	Coverage (Recall)	Efficiency By Chance	Coverage By Chance
10	1,510	20,207	5,025	67,855	23.1%	7%	23%	7.1%
9	3,148	18,569	10,405	62,475	23.2%	14.5%	23%	14.7%
8	3,228	18,489	10,736	62,144	23.1%	14.9%	23%	15.1%
7	11,562	10,155	25,180	47,700	31.5%	53.2%	23%	39.8%
6	14,320	7,397	34,715	38,165	29.2%	65.9%	23%	53.2%
5	17,547	4,170	49,753	23,127	26.1%	80.8%	23%	73%

Figure 1: Results for prioritization strategies based on CVSS base scores

Let's say an organization sets a policy of remediating all vulnerabilities with a CVSS score of 8 or more (CVSS scores range from 1 to 10, so that's generally considered high severity). They would achieve coverage of just 14.9%. Randomly patching "by chance" (rightmost column) yields essentially the same result of addressing just 15.1% of exploited vulnerabilities. I don't know about you, but that definitely exceeds my risk tolerance.

> ***Remediation Coverage:*** *Measures the completeness of remediation efforts—that is, the percentage of exploited or "high-risk" vulnerabilities that were actually remediated.*

A more risk-averse approach would be to lower our remediation threshold to CVSS 5 or higher. That has the positive effect of boosting coverage to 81%, though false positives (vulnerabilities remediated unnecessarily because they were never exploited) skyrocket nearly fivefold. So, while my VM team is indeed fixing most of the riskiest bugs, we're wasting a ton of time and effort doing so, and we're still not performing much better than random chance.

REMEDIATION EFFICIENCY

The concept of wasting time remediating the wrong (less risky) vulnerabilities segues right into another metric I wish was more common among VM programs—remediation efficiency. This metric measures the precision of remediation efforts—the percentage of vulnerabilities remediated that have actually been exploited.

I like efficiency because it acts as a counterbalance to coverage. As seen in Figure 1 from the previous section, I can drive up coverage just by fixing more vulnerabilities. But that comes at a cost and quickly becomes intractable. Adding efficiency into the metric mix can help teams identify when their remediation efforts aren't achieving an acceptable ROI.

> ***Remediation Efficiency:*** *Measures the precision of remediation efforts—i.e., the percentage of vulnerabilities remediated that have actually been exploited.*

It's possible to measure just one or the other, but coverage and efficiency are metrics that work best in tandem. Risk-averse organizations may prioritize coverage, thereby remediating more than they truly need to. Resource-constrained programs may start with flaws actively being exploited in the wild to maximize efficiency. Others may attempt to strike a balance between them. But the important point is that both metrics provide useful, actionable insight into the performance of the VM program.

Wondering where organizations stand when it comes to coverage and efficiency? Well, here's the chart for you. We tracked this metric for hundreds of organizations over several years. Each dot in Figure 2 represents an organization and its performance for coverage and efficiency. Notice that it's rare for organizations to rate high on both scales, indicating there's often a tradeoff between these metrics based on organizational strategies and capabilities.

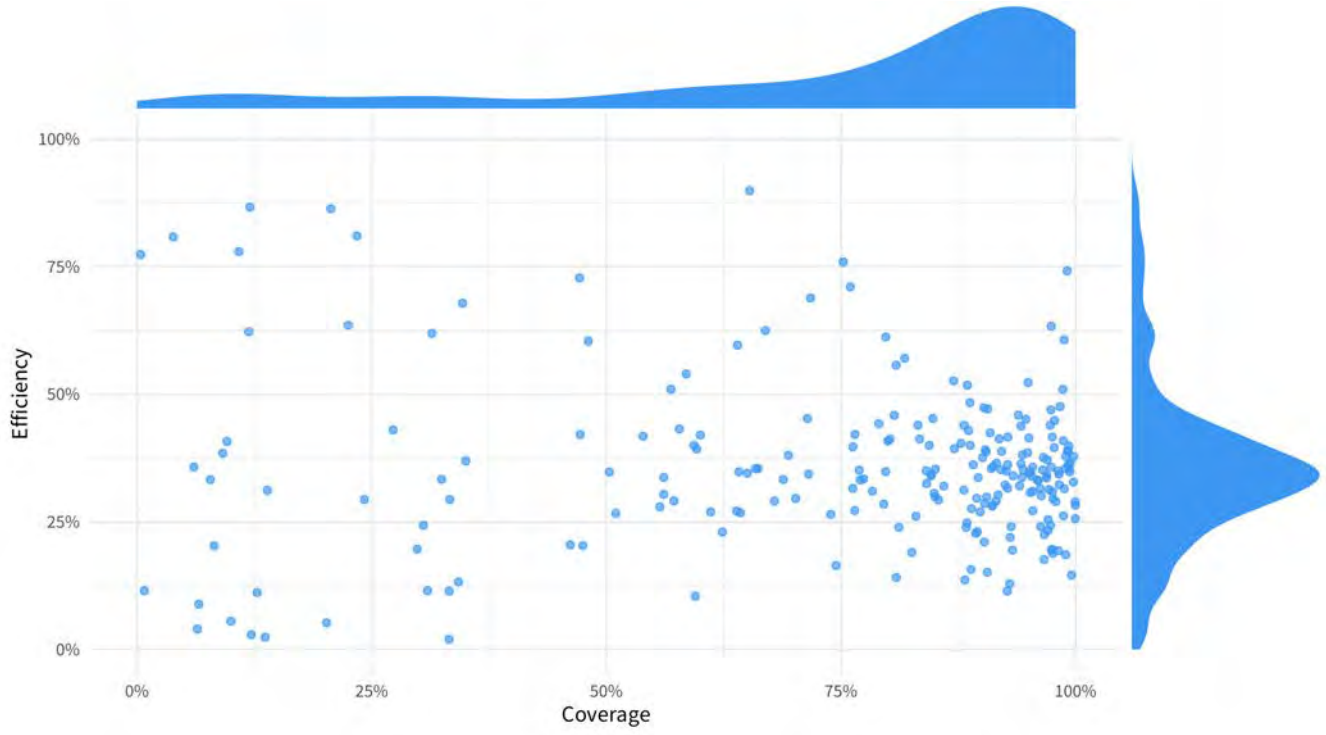

Figure 2: Remediation coverage and efficiency measured in several hundred organizations

REMEDIATION VELOCITY

If an organization successfully achieves 100% coverage (remediates all exploited vulnerabilities in their environment) but takes forever to do so, it will completely undermine that achievement. Once a vulnerability is published and patches are available, there's a race against time to remediate them before affected systems are exploited. What does that timeline generally look like? Figure 3 lays it out well.

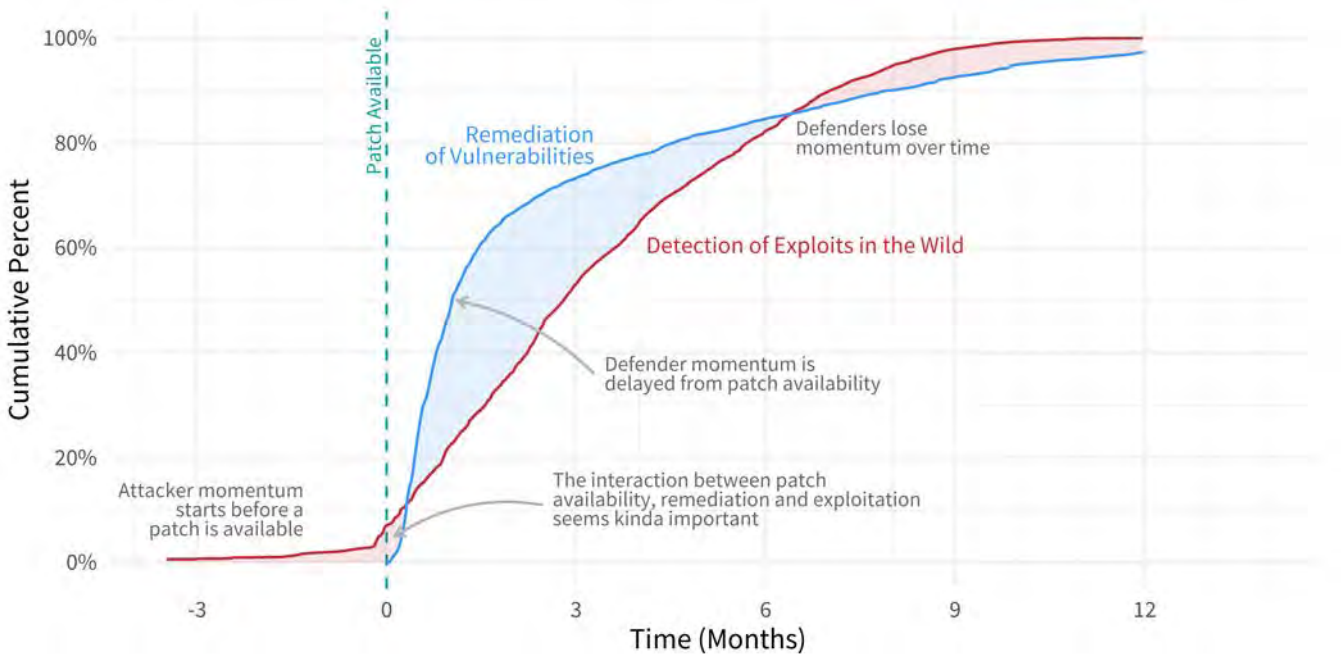

Figure 3: Comparison of vulnerability remediation and exploitation timelines

It generally takes VM teams about five months to remediate 80% of vulnerabilities across their environment. Attackers typically reach 80% of their target population with exploit attempts in about six months. Who's winning is often a matter of what part of the race you check. In the time leading up to a vulnerability being patched, attackers have a solid advantage. Armed with the patch, defenders take back the momentum by remediating exposed vulnerabilities quickly. After that initial push, defender momentum wanes, and attackers regain the long-term advantage. Fascinating, right?

Most organizations can't measure the rate of exploitation activity spreading in the wild (though their service providers might be able to help), but they absolutely can and should measure their own remediation velocity. There are several ways to measure that, three of which are included in Figure 4, which shows a "survival curve" for vulnerability remediation across hundreds of organizations.

> ***Remediation Velocity:*** *Measures the speed and progress of remediation, or how quickly issues are addressed and how long they persist within and/or across assets.*

Survival analysis is a classic technique for measuring the time duration for an event. For our purposes, the event of interest is remediating a vulnerability, and it's a very useful way to study how long this remediation process takes.

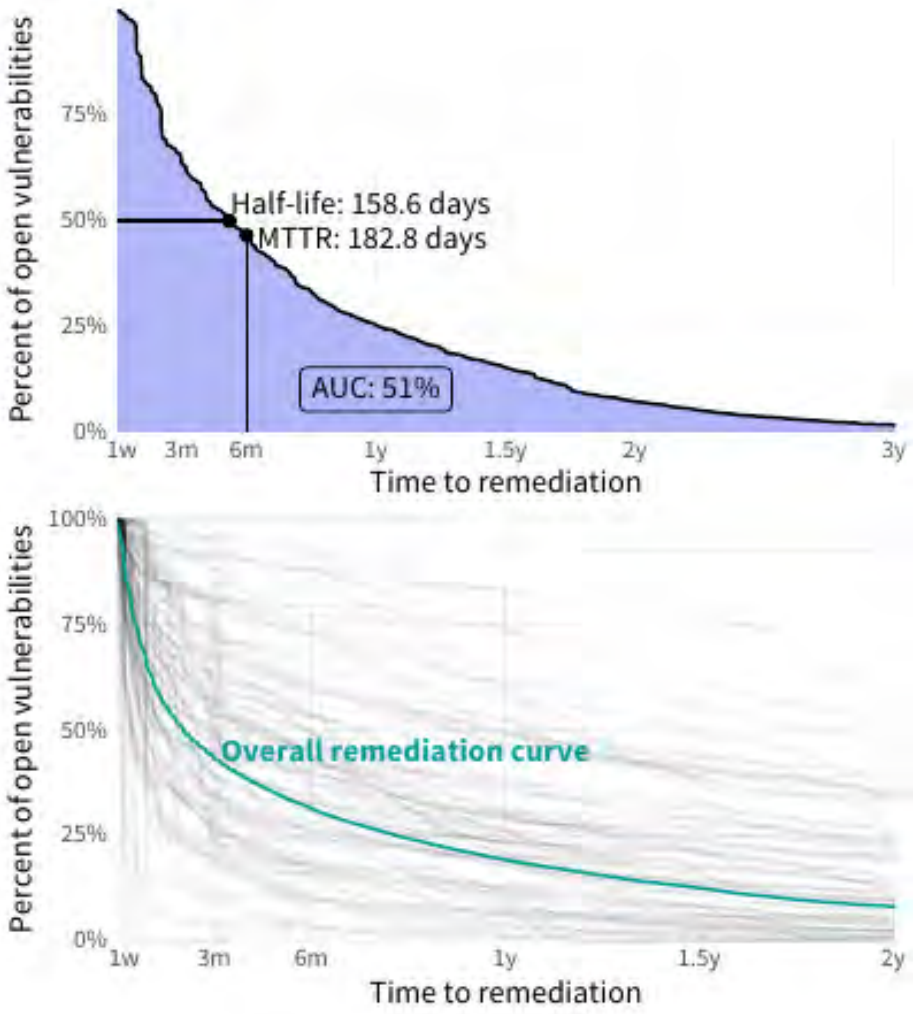

Area Under the Curve (AUC): Measures the area under the survival curve representing "live" (open) vulnerabilities. A lower AUC means higher velocity.

Mean Time To Remediation (MTTR): Measures the average amount of time it takes to remediate vulnerabilities.

Vulnerability Half-Life: Measures the time required to remediate exactly 50% of open vulnerabilities.

Figure 4: Survival curve for vulnerability remediation velocity and associated metrics

Note that the vulnerability survival curve (or remediation timeline, if you prefer) shown in Figure 4 is an overall view. Remediation velocity will vary across different types of organizations, assets and vulnerabilities. And you can apply these remediation velocity measures to all vulnerabilities, or just to those with known exploits, or to both. If you wind up tracking both, I advise setting more aggressive remediation timelines for the latter. You'll see why in the next and final section.

VULNERABILITY EXPLOITATION

I know I said "top 3 metrics," but you probably noticed that the three discussed above require knowledge of whether or not vulnerabilities have been exploited. That is indeed correct, and that's why I've included the tracking of exploits as another critical thing VM programs should do if they want to create meaningful metrics that drive performance and risk reduction.

By "exploitation," I refer to either published exploit code or active attacks against organizations. I view this as critical intel for VM teams, because exploited vulnerabilities are the ones that represent a real and present danger for organizations. Exploit code acts as an early warning sign because, as seen in Figure 5, exploitation activity in the wild jumps dramatically once exploit code publishes.

Exploit code: *When a vulnerability becomes public, proof-of-concept or working code for exploiting it may be published as well. Exploits can be traded in the digital underground, shared on above-ground mailing lists, published in GitHub or included in exploitation frameworks like Metasploit.*

Active Exploitation: *Once threat actors (or red teams) begin probing for and attacking vulnerabilities in organizational assets, this constitutes active exploitation "in the wild." This activity can be observed through security technologies (e.g., IPS), malware analysis, etc.*

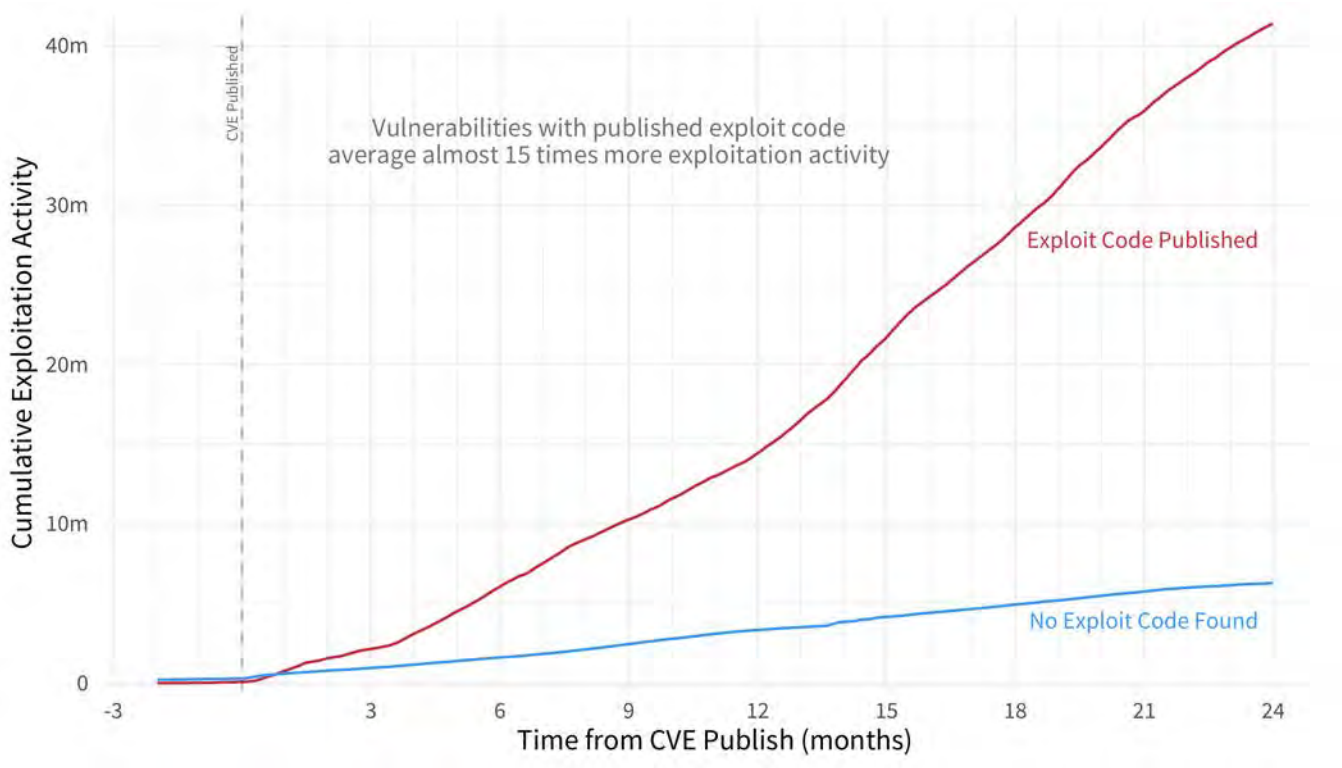

Figure 5: Exploitation activity detected for vulnerabilities with vs. without published exploit code

Most metrics are based on things measured by the organization itself, but this one relies on intel from the outside. That begs the question of where to get intel around the publication of exploit code and active attacks in the wild to use in tracking this metric. There are lots of companies that sell services for this, but I'm not going to get into that here. Instead, I'll share some free sources you can begin monitoring right now.

First, the Cybersecurity and Infrastructure Security Agency (CISA) maintains a **catalogue of Known Exploited Vulnerabilities** (KEV), which is exactly what the name implies. Federal agencies are required to remediate all vulns in that catalogue within a certain timeframe. Savvy private sector companies could begin measuring remediation coverage against the KEV too.

I already mentioned Metasploit as a source for tracking exploit code, but GitHub is another. In fact, **Cyentia Institute research** shows that GitHub has overtaken legacy exploit sources like Metasploit and Exploit DB as the primary place for publishing exploit code.

Finally, I recommend familiarizing yourself with the **Exploit Prediction Scoring System** (EPSS). EPSS is a community-based and data-driven effort that estimates the likelihood that a vulnerability will be

exploited in the wild. There's a lot of momentum around EPSS lately and, best of all, it's free! Ditch the CVSS-based metric you're using now to prioritize vulnerability remediation and replace it with an EPSS-based metric. Here's some **analysis that justifies why.**

In addition to focusing remediation efforts on the riskiest vulnerabilities (coverage), tracking exploits can dramatically reduce workload (efficiency). Let's walk you through one last chart that demonstrates how that works.

Figure 6 separates vulnerabilities into those actually observed in a live environment and those that were never seen. On the right, two-thirds of published CVEs were never detected by the hundreds of firms in our sample. That suggests many vulnerabilities affect technologies not currently used in most enterprise networks.

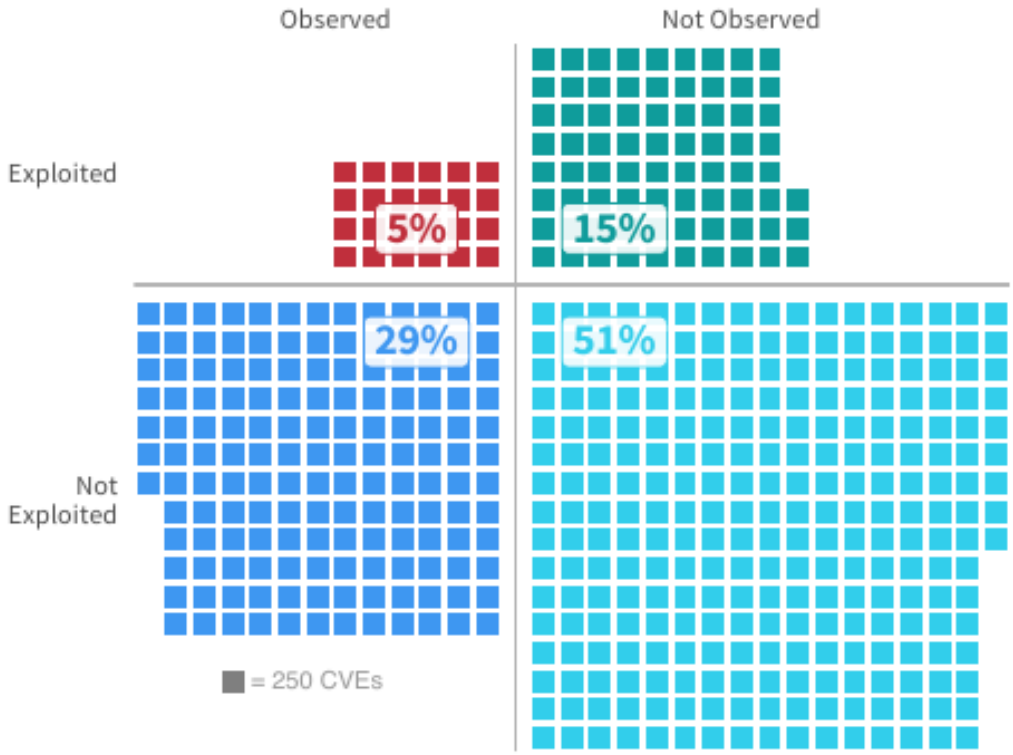

Figure 6: Ratio of observed/not observed and exploited/not exploited vulnerabilities

We then separate CVEs into those that have been exploited and those with no known exploits. Even though 20% of published CVEs have a clear threat (either actively exploited in the wild or a published exploit exists), only about 5% of them exploit software or hardware commonly used by organizations. While it's certainly true that new exploits emerge regularly, it's also true that most firms can safely deprioritize vulnerabilities without known exploits.

THE BOTTOM LINE

When you can't fix everything, fixing what matters most is critical. And that's where I see hope for beleaguered VM programs. The research highlighted in this article definitively shows that organizations able to maintain this risk-based focus remediate vulnerabilities better, faster and more efficiently than those with less-focused programs. I hope these three metrics help your VM team develop and mature that risk-based focus!

*Note: All charts in this article are taken from the **Prioritization to Prediction** series of reports jointly produced by Kenna Security and the Cyentia Institute.*

WHAT CYBERSECURITY METRICS MEAN UP AND DOWN A COMPANY

DAVID HECHLER

Andy Geisse has seen the problem again and again at start-ups. Executives are focused on building the business. They're concerned about metrics, but the ones that have their attention involve sales and metrics like KPIs. For them, cybersecurity is an "afterthought," he said. In some cases the companies haven't even identified a framework against which to measure their cybersecurity risks.

Geisse (rhymes with ice) is an operating partner at **Bessemer Venture Partners,** which invests in quite a few new companies—and analyzes many more. During a long and distinguished career that preceded his current position, Geisse worked at various companies—sometimes as a CIO, sometimes as a CEO—where he gathered and passed along data, and made business decisions based on different sets of metrics. Given this background, he's been asked to join the boards of start-ups, where he's seen the deficiencies up close.

He offered an example. In 2015, he said, after he'd joined the board of a company he declined to name, he kept bringing up cybersecurity. Far from poring over the latest data, no one there seemed to be paying close attention. "It's done by the audit committee," he was told. When he spoke to the committee, "it was clear they were mainly CFOs, or ex-CFOs, who really didn't understand cybersecurity."

He'd seen this before. "Often engineering companies" like this one "think they know how to do it all," he explained, "and don't understand the complexity and importance of IT, and having somebody focused on that area." After a lot of pushing, the company appointed its first CISO and its first CIO.

Geisse urged the new hires to conduct an audit, and they brought in an outside expert to help. "It was shocking the amount of gaps we had," he said. Not long after,

Andy Geisse

"I think it's absolutely critical that boards start taking [cybersecurity] as a focused issue in the governance of the company."

they discovered that somebody in marketing had pulled customer data and stuck it in an AWS database with no security around it. "That caused us a hit in terms of notifying those customers," he said.

The lessons? "I think it's absolutely critical that boards start taking this as a focused issue in the governance of the company," he said. "When a board has somebody that has some cybersecurity background," he continued, "it can make a huge difference in reducing those risks."

THE OLD DAYS

Cybersecurity wasn't always so complicated. When he first served as a CIO, there were no cloud services. "Almost all your applications were behind your firewalls in your data centers." A lot of the metrics he needed "I would carry around on a one-page document that I would keep with me at all times." Many concerned "projects and development capacity and how we were doing against business goals."

"Every time there is a major hack that hurts a company, hurts their reputation, all of a sudden boards are asking the question: 'What are we doing about it?'"

Later, when he was a CEO in the early 1990s running wireless companies, "it was more about sales and financial metrics, in terms of running the company," he noted. "And frankly you expected your CIO to handle the systems and handle cybersecurity. You never thought about it." And for good reason, he added. "In those days a lot of the systems weren't even connected to the outside world."

That had changed by the time he became CEO of AT&T Business Solutions in 2012. He was no longer cruising at 25,000 feet, above and beyond the world of IT. "I would often have to report to my board issues we had that potentially exposed us to lost customer data or lost employee data," he recalled. "I still depended on my CIO and my CISO to run and manage those areas, and have their own metrics, but I also expected regular reports from them on how we stood." He needed to know about hacks, and how prepared the firm was—and how quickly the damage was being repaired.

THESE DAYS

Today, boards are attuned to potential risks to their brands. And cybersecurity is on their watch lists. "Every time there is a major hack that hurts a company, hurts their reputation," Geisse said, "all of a sudden boards are asking the question: 'What are we doing about it?' And: 'Are we protected?'" The attention ramps up for a time. Until the news cycle passes and fear is replaced by another long lull. Which is only broken by the next headline hack.

Another pattern Geisse has noted is the plethora of cybersecurity solutions. "All of a sudden there's a million tools out there," he said. It's not necessarily a good trend, he added. "I think it's much more about having the right framework, and then auditing against that framework to make sure you're protected." There are plenty of frameworks, he said, but he prefers traditional metrics: the number of hacks, the location of intrusions, the response to them, the speed of mitigation.

But one area that demands special attention, he said, is data. There are many more metrics on data for the simple reason that there's more of it, and it's everywhere. "Your data is no longer just in your data centers," he observed. "It's on Google. It's on AWS. It's on Azure. It's on SaaS-based systems." So a big trend, he said, is understanding where it is, and making sure that it's protected in all those places.

> "I've seen more headhunters looking for board members with a cybersecurity background. I know that just from the folks that have called me."

This reminded him of another example. It involved a large retailer that had "pretty good cybersecurity," Geisse said. All the "normal things" looked good. But customer complaints had told them that something was wrong. It took them many months to discover the source. Hackers had inserted a small piece of software in their point of sale systems, which proved to be the origin of a data leak. The hackers exfiltrated credit card numbers, but slowly—slow enough, he explained, that they wouldn't show up in an average audit. That made the hack very hard to detect, like an automobile tire with a slow leak. Ultimately they analyzed IP addresses "and found one that shouldn't have been there."

THE ROLE OF THE BOARD

While Geisse believes that boards of directors benefit when they have members who understand cybersecurity, he doesn't think they want to be briefed on the company's cyber metrics. They don't need detailed lectures from the CISO or CIO. But he does think an important step can be having a board member with this expertise sitting on the board's audit committee or a cybersecurity committee. Those committees can communicate directly with the CISO and the CIO, and report relevant information back to the full board.

It's likely that the board will continue to respond to headlines that remind them of the threat to the company's reputation, then drift back into a lull of complacency. But having regular access to more and deeper information can only help make the company safer, he said.

Asked if he thinks this is something that's catching on, he nodded emphatically. "In the past five years," he said, "I've seen more headhunters looking for board members with a cybersecurity background. I know that just from the folks that have called me. So I do think there's a positive trend there."

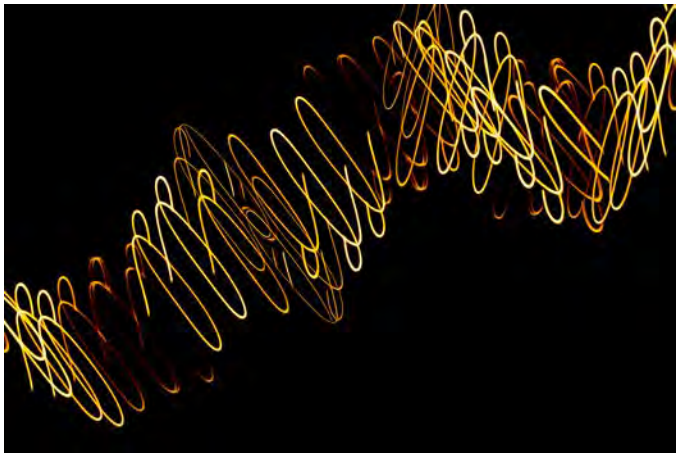

SIX QUESTIONS TO ASK WHEN EVALUATING CYBER THREAT INTELLIGENCE PLATFORMS

CHRISTOPHER R. WILDER

Cyber threat intelligence (CTI) is one of the fastest growing areas of cybersecurity, especially among security operations and security operations center (SOC) teams. CTI information is how the organization gains visibility into potential threats to its infrastructure, assets and employees. Increasingly, it seems pivotal. And yet, I had a conversation with a colleague recently that stopped me in my tracks—and made me realize that sometimes there's a hole in this approach.

"CTI platforms are great," said my colleague, Dr. Jennifer Bayuk, "but how do users measure and ensure the efficacy, timeliness, actionability and accuracy of the data they consume?" Good question!

There is no mechanism for benchmarking how enterprises ingest CTI data. Tactically, CISOs and SecOps teams focus on controls and systems to block and revoke network access in case of a compromise or breach. Strategically, most lack the skills or personnel to understand behavior analytics and manage disinformation campaigns or incoming threats.

Security teams commonly focus on threat indicators, also known as Indicators of Compromise (IoCs) and tactics, techniques and procedures (TTPs) as the foundation of their CTI programs. While this is a good start, every attack vector must be addressed when facing an expanding bad actor landscape. For example, tactical threat indicators are data that associate observed artifacts such as website URLs, file hashes, DNS/command and control (C2), or IP addresses with known threat activity such as phishing, botnets, ransomware and malware threats.

Every business knows that an attack and eventual compromise will occur. How the organization responds makes all the difference.

MEASURING AND BENCHMARKING THE EFFECTIVENESS OF CTI DATA FEEDS

Every business knows that an attack and eventual compromise will occur. How the organization responds makes all the difference. CTI plays an important role in any organization's defense strategy. It's leveraged and integrated into other cybersecurity functions, such as XDR, SIEM systems, incident response and forensic investigations. So, what can enterprises do to ensure that they have timely, relevant and accurate CTI information?

We believe CTI feeds must align with the organization's threat models, security team capabilities and ability to respond to attacks. Based on that premise, here are six key questions to ask a CTI vendor when you select an intelligence provider—and why they are important for you to understand.

We anticipate over 70% of security operations teams leverage CTI information to identify, react to and block security threats as they occur.

1. IS THE CTI DATA ALIGNED WITH MY ORGANIZATION'S THREAT MODEL AND SECURITY GOALS?

Commonalities exist in every organization, regardless of industry or sector. However, no two businesses are exactly alike; there is a high likelihood that each one will have a different risk tolerance or attack vectors. Because of these factors, enterprises must deploy CTI solutions that improve visibility into their threat models, deliver both contextual and actionable intelligence, and help SecOps to react to compromises and breaches faster.

Threat modeling and planning are critical to ensure that resources, controls and real-time intelligence address actionable threats to the organization, employees and infrastructure. Therefore, to ensure sourced CTI is effective, it must align with an existing threat construct.

For security operations teams, it is important to acquire intelligence sources that align with and provide visibility into the organization's threat vulnerabilities and desired outcomes. Security teams must benchmark and incorporate intelligence data that addresses their needs.

2. WHY IS IT IMPORTANT TO ENSURE THE ACCURACY, QUALITY, RELIABILITY AND CONSISTENCY OF CTI?

One of the challenges most SecOps teams have is the lack of skills or tradecraft to turn raw data into contextual information and make it actionable intelligence. CTI and vulnerability information originates from multiple internal and external resources. In many cases, CTI data is ingested and processed using automation. Because teams lack the skills to identify real-time threats, many threats get past SecOps and cause havoc within the company.

Whether derived internally or externally, CTI quality, reliability and consistency are critical to improve the organization's cybersecurity posture. Internal sources include proactive vulnerability scanning, intrusion detection and endpoint detection and response systems, network monitoring and behavioral analysis tools. External CTI sources come from threat intelligence platform vendors, commercial and community collaboration forums and other intelligence feeds.

Further, improving the accuracy of CTI feeds helps organizations reduce false positives in an intelligence report. False positives imply poor data quality, forcing SecOps teams to pursue "fool's errands" and investigations that lead nowhere and put their companies at risk.

3. HOW DO I BALANCE THE COMPLETENESS AND RELEVANCY OF CTI INFORMATION?

Visibility of threat information gives enterprise teams a more complete understanding of how best to react and respond to cybersecurity threats. An accurate and relevant intelligence feed allows SecOps and SOC teams to ask the right questions and act accordingly. Having well-established threat models and a proactive attacker point of view enables SecOps teams to make more informed decisions when responding to a nation-state or determined hacker organization.

Effective CTI provides actionable insight into the bad actor, their capabilities, targets and attack vectors.

Bottom line: The more your organization and industry context are available within CTI, the more useful it is. Moreover, internal or organic threat intelligence should have more weight than externally sourced intelligence information, as it is more difficult to take external and generic data feeds and make them contextual.

4. WHAT ARE THE TRAITS THAT MAKE CTI INFORMATION USEFUL AND RELEVANT?

We anticipate over 70% of security operations teams leverage CTI information to identify, react to and block security threats as they occur. Further, nearly 40% of security operations proactively use CTI to deter active threats. For SecOps teams, the top use case for CTI in security operations is to use CTI to detect and prevent potential cybersecurity events while locating sources and methods to eliminate malicious activities or threats. When considering a CTI feed, this is a key criterion that must be part of the decision process.

An effective security monitoring strategy correlates and analyzes data from multiple sources to detect threats before they can cause harm to the organization. Reliable and consistent CTI sources are one way to ensure the optimal use of intelligence on resources and assets, allowing SecOps teams and SOCs to focus on threats that pose the highest risk.

Finally, continually reviewing security monitoring procedures is critical to determining how much CTI influences monitoring strategies and which sources are right for your enterprise.

5. WHAT IS THE IMPORTANCE OF INTEGRATING CTI WITH INCIDENT RESPONSE PROGRAMS?

Visibility and context into threats, compromises and breaches allow enterprises to respond to security incidents adequately. Effective CTI provides actionable insight into the bad actor, their capabilities, targets and attack vectors. Having visibility into the motivations, TTP and the anatomy of an attack gives security teams assurance that they can deploy appropriate defense mechanisms to prevent a successful attack from determined hackers and crackers.

As part of your review, it is critically important to assess how CTI information is integrated with the steps in your organization's incident response approach, including preparation, detection, analysis,

containment, eradication and recovery. Further, incident response teams should have enough insight and information from their CTI feeds to influence and guide the development of contextually relevant cybersecurity playbooks and flight guides.

6. WHAT IS THE IMPACT OF HAVING A PROACTIVE CTI PROGRAM?

Most enterprise teams perform a postmortem review only after a breach, compromise or security incident. Being proactive allows teams to understand inbound threats and gain insight into what worked well (and what did not) during incident detection. It also helps them enhance their responses and identify improvement opportunities. Unfortunately, most security organizations do not leverage CTI as a proactive measure.

Security teams must have clear visibility and context into security incidents to determine whether their CTI feeds and monitoring technologies can respond to and detect unknown and zero day threats. Security teams understand that reducing time to identify and respond to threats, and preventing significant damage to systems and data, help organizations mitigate disruptions and improve their overall security posture.

Assessing relevant CTI data reduces the impact of security incidents. It provides insight into which intelligence sources deliver the best value to the organization, and it justifies continued investment in security teams.

THE TAKEAWAY

The value of CTI to any organization is the ability to support timely, contextual and actionable threat information to enterprise stakeholders at all levels of the organization, from the board and executives to CISOs, SOCs and SecOps teams. Those teams must choose their CTI provider carefully, ensuring that the information they ingest is relevant and actionable for their enterprise. Quite simply, asking the right questions and evaluating CTI sources help SecOps teams detect and mitigate threats.

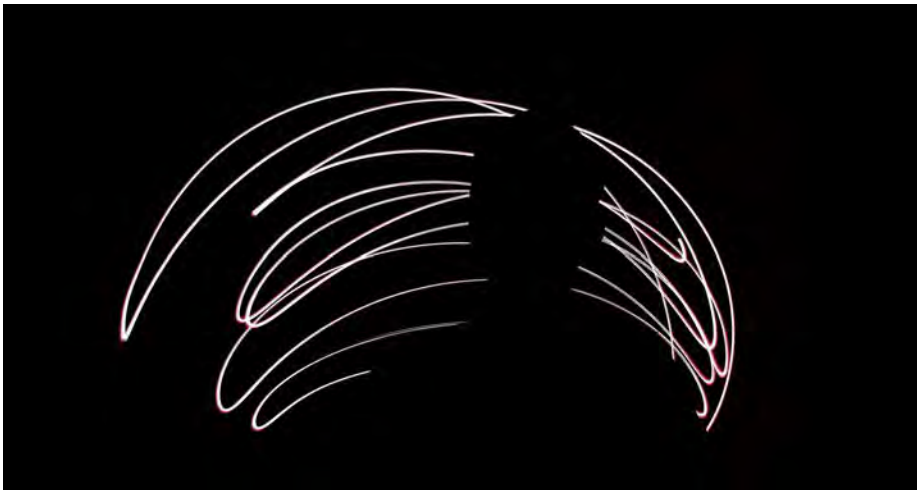

THREE LESSONS FROM THE HISTORY OF CYBERSECURITY METRICS

DR. EDWARD AMOROSO

My colleague Jennifer Bayuk recently took on the task of chronicling the history of cybersecurity metrics. She did so in a short video and associated technical report—and I recommend both to any enterprise security practitioner or academic.

As you probably know, our security industry has become obsessed with metrics, despite the fact that many participants can barely define what they mean versus, say, a measurement or a scale. And the result is that we have a lot of bad metrics, as you also probably know.

As a short teaser to Jennifer's fine works, here are three key lessons that I learned from them. I hope my summary helps to underscore the importance of reviewing our past to improve our present and future.

1. METRICS THAT SLOW DOWN BUSINESS PROCESSES WILL GO AWAY

In the 1980s, the now-defunct U.S. National Computer Security Center created a set of criteria for establishing and comparing how well an operating system has implemented security. Its work resulted in the so-called Orange Book (aka DoD 5200-28-STD).

The way it worked was that you toiled like crazy to add security control commensurate with some targeted criteria class. Then a team of external assessors would come in and review what you did, resulting in a certification and report. It all took many, many months.

Jennifer's account brought back nice (and also tough) memories for me, because I personally spent many, many years working as part of a group that built a Unix team certified at the Orange Book B1 level. Our biggest customer was Harris Corporation, which had constructed a firewall on our OS.

The Tower of Babel was a perfect metaphor for cybersecurity researchers trying to build a taxonomy without agreeing on terminology.

But the Orange Book ultimately failed. Why? Because of that many, many months (and many, many years) thing. This was all before DevOps, but certification definitely slowed things down. Once certified, you literally couldn't touch the software—and this would never work.

So, we learned a valuable lesson from the Orange Book era: Metrics that slow down business processes will go away. This was the case with both the 5200-28 criteria and even the entire NCSC, which was quietly dissolved in the '90s.

2. KEEPING TRACK OF VULNERABILITIES IS TOUGHER THAN IT LOOKS

The Tower of Babel is a mythical structure mentioned in passing in the Bible (Gen 11:1-9). It was a huge tower built by humans to reach up into the heavens, but workers eventually gave up and the structure was never built.

The story is prominent because it is used to explain why humans speak so many languages. Interpreters claim that God decided to dissolve human linguistic ability into many families of language, not to mention different, isolated groups with different features.

In the '90s, computer security practitioners used this mythical story to explain how vulnerability research was progressing. In particular, what was happening was that different groups were using different language, references and definitions to describe the same software bug.

This was the first hint—one since confirmed—that it was going to be difficult to build useful vulnerability taxonomies if vendors failed to coordinate. Efforts such as the National Vulnerability Database were started, but they did not solve this issue.

So, we learned an important lesson from the Vulnerability Tower of Babel. Keeping track of vulnerabilities is tougher than it looks. And despite many attempts to improve coordination and standardization, this remains a problem.

3. SECURITY METRICS STILL HAVE A WAY TO GO

Jennifer's historical report demonstrates something interesting and also just a tad humorous. Specifically, she points out that as the complexity and challenges associated with metrics have grown, we've reached the ultimate irony: metrics for metrics.

When you review Jennifer's history of our attempts to use metrics to explain our discipline, it should come as no surprise. With so many different measurements, scales and metrics for cybersecurity, it's natural that confusion would arise.

And so, to explain the Tower-of-Babel metrics delivered by SIEMs, SOARs, EDRs, GRCs and the like, we go into the boardroom and subject executives to metrics that explain our metrics. Witness the many commercial tools pitching metrics dashboards.

The only reasonable conclusion that can be drawn from this situation is this: Security metrics still have a way to go. Let's hope that the future focuses on simplicity and elegance because otherwise, like the builders of that tower, we might be forced to give up and go home.

A BRIEF HISTORY OF CYBERSECURITY METRICS

DR. JENNIFER BAYUK

Measuring security is hard. It is hard because security is ephemeral. Like the weather, you don't measure security directly, you measure other things then try to make predictions about the environment. Weather people measure temperature, pressure and wind. Cybersecurity people measure threats, vulnerabilities and the effectiveness of controls. Both try to predict the stability, or lack thereof, in the target environment.

What many people don't know about measuring security is that it is even possible. Tools and techniques for security measurement are just as plentiful as tools and techniques for studying the weather. That does not mean that they are extremely accurate, any more than weather equipment is 100% effective at predicting rain within a 30-square mile area on a given summer day. But just as the National Weather Service is getting pretty good at pinpointing the likelihood of tornados within city limits, cybersecurity professionals are getting good at predicting whether a company will be the target of an attack, and if so, what impact to expect.

In "The Coming Storm," Michael Lewis reviewed the history of weather forecasting and noted that the first professional weather forecasters went into business during World War I not because the capability for accurate forecasting had emerged, but because the need for forecasting was so great. He quipped that the weather forecast supply industry had to commit fraud to keep up with demand.

In the 1960s, computer security metrics were introduced to business professionals because mainframes were used to calculate financial statements.

Today, the same could be said of some, though not all, cybersecurity metrics. Here we will shed light on the situation by reviewing some of the mileposts that have been successfully passed over the last 60 years.

Before the 1960s, military and secret government projects had studied security metrics even before the ubiquity of computers (they measured cryptographic strength by how long it takes to decrypt an encrypted message). In the 1960s, computer security metrics were introduced to business professionals because mainframes were used to calculate financial statements. The accounting profession required some way to measure the integrity of the data to show that financial statements were correct. They created the Electronic Data Processing (EDP) Auditors Association, a group of auditors that came up with an input-output test for computer integrity, called "auditing around the computer."

People used to keep manual records of everything they typed into the computer, and the auditors would collect the manual records, test a statistically valid sample of them with manual calculation, and if they got the same result as the computer, they declared the computer processes to be accurate. We still have tests like that. We don't do manual calculations, but we do test computers against each other. And we call it an integrity metric. The auditors started a list of ways to ensure computers archived business objectives, called "control objectives," just as they would in any management environment. We still use lists of those to develop cybersecurity metrics (see Figure 1 below).

Figure 1: Control Objectives

With the advent of personal and mini-computing in the 1970s and 1980s, it became harder to test all the programs that were on a computer. And it was beyond the capability of any one technology auditor to come up with control objectives to cover the exploding number of threats to computer data and operational integrity, much less to measure them. Auditors found kindred spirits in a government-sponsored, public-private group that was charged with creating standardized terminology for computing in general, and trusted computer system evaluation was a very significant component of that effort. This group collaboratively developed standards in a manner similar to the current National Institute of Standards and Technology (NIST).

The U.S. government was really concerned about military computers—that is, making sure secrets were safe. And so the standards body came up with ways to describe what safety meant by creating symbols for computer subjects (i.e. users) and for computer objects (e.g. files), and creating logical theorems around how subjects could access these objects, whether running them or reading them or writing to them. And through that work, they came up with basic access control principles we still use today, and ways to measure that subjects could be, and in fact were, limited to a minimal set of objects.

The computing standards were set out in a publication called the "Rainbow Series"—so called because the books were identified by color. Security measures were in **"The Orange Book,"** which was given out by the government and industry associations to security professionals (see Figure 2 below). It was through this book that many of the cybersecurity pioneers learned that measuring security was possible.

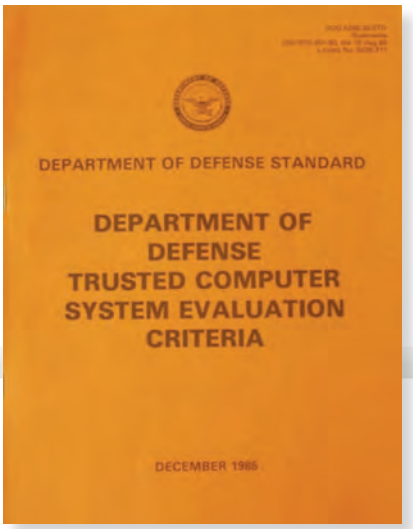

Figure 2: "The Orange Book"

In the 1980s, the computing market exploded with alternative subject-object implementations on diverse computer operating systems. Both local area networks (LANs) and business-to-business computer services offered telecommunications networks ("dial-up" services, now known as "third parties") that became ubiquitous. This made it impossible for the theoretical demonstrations like those in "The Orange Book" to keep up. So another consortium was formed to establish what was "common" about all those diverse technologies that could be codified in a set of instructions about how to plan and implement your own unique "protection profile." This profile is the combination of technologies that provide assurance that your business security objectives are met on all of your computing platforms.

The "Common Criteria" provided a methodology to develop and measure our security targets (also known as requirements), as well as to formally verify that our security protection profile features perform as expected via methodical testing. The logic is that the more diligently you follow the methodology, the more assurance you have that your security works (see Figure 3 below). A plethora of such methodology standards followed for both security implementation and assurance. These metrics provided at least some evidence to auditors and senior management that computers were under control.

As we approached the bicentennial, the information security marketplace exploded with products designed to help information security officers provide such assurance. Antivirus was the first, followed by operating system security and log consolidation vendors.

Unfortunately, the vendors all seemed to be anti-standard, each claiming that their product was unique and thwarting efforts by industry analysts to put them in well-defined categories. Antivirus vendors gave different names to the same virus and each claimed to have discovered it first. A famous paper

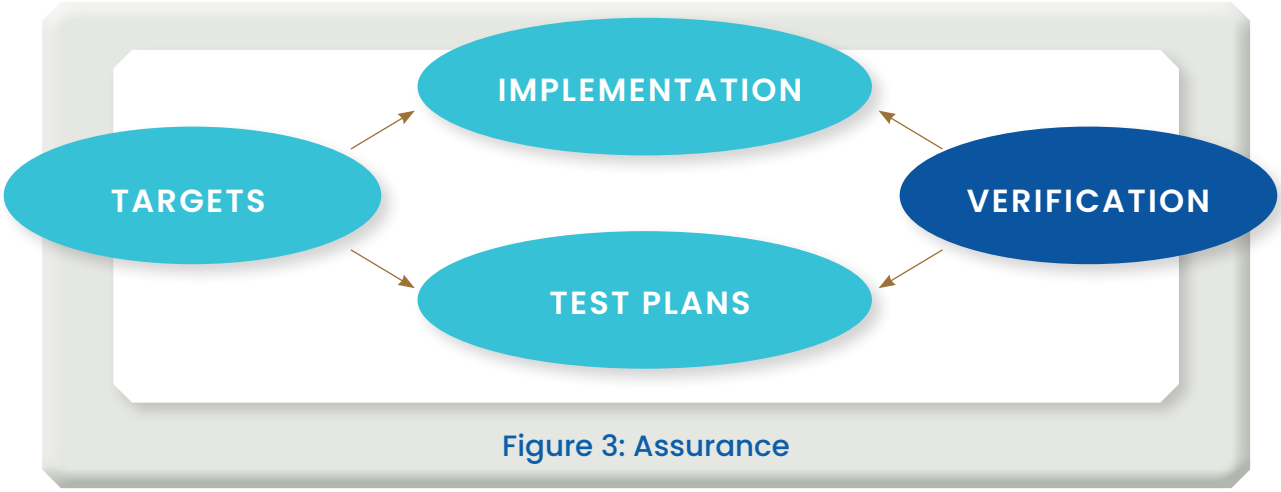

Figure 3: Assurance

called this situation the **"Vulnerability Tower of Babel"** (see Figure 4 below). In response, DARPA and NIST sponsored a conference to discuss setting standards for vulnerability labeling. The effort was merged with the Common Vulnerabilities and Exposures initiative started by MITRE around the same time. Ongoing efforts flowing from this initiative led to the creation of the National Vulnerability Database (NVD) and the Common Vulnerability Scoring System (CVSS).

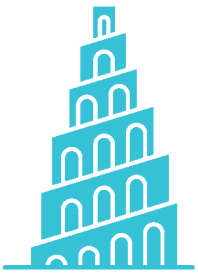

Vulnerability Tower of Babel
-1998-

Organization	Name Referring to Vulnerability
AXENT *(Now Symantec)*	phf CGI allows remote execution
BindView	#107 — cgi-phf
Bugtraq	PHF Attacks — fun and games for the whole family
CERIUS	http_escshellcmd
CERT/CC	CA-96.06.cgi_example_code
Cisco Systems	HTTP — cgi-ph
CyberSafe	Network: HTTP 'phf' attack
IBM ERS	ERS-SVA-EO1-1996:002.1
ISS	http — cgi-phf
Symantec	#180 HTTP server CGI example code compromises http server
Security Focus	#629 — phf Remote Command Execution Vulnerability

Figure 4: Vulnerability Tower of Babel

In the early 2000s, data related to information security was starting to pile up. Best practices in managing security included review of logs of unauthorized access attempts, authorized access at unusual hours and server process execution outside of system boot sequence. Wherever computer operating systems had requirements for confidentiality, integrity and availability, logs were generated. Where the operating systems could not implement the requirements, security software was "bolted on," and the security software logs had to be reviewed as well.

In time it became impossible to review all the logs without automation. Security Incident and Event Management (SIEM) technology allowed us to put logs in a database and receive an alert if they contained patterns that appeared suspicious. Large budgets for Security Operations Centers (SOC) were becoming the norm. A periodic meeting of specialists in cybersecurity metrics was arranged by industry analysts, drawing dozens of volunteer program committee participants as well as sponsors. They called it **Metricon**.

Between 2004 and 2019, Metricon produced detailed definitions for dozens of security metrics categories. For example, in 2009 Jim Cowie from Renesys shared a new method for monitoring internet routing vulnerabilities, downtime and instability, hijacking and wholesale traffic interception. **His presentation** is an exemplar in that it provided a concrete way to measure an extremely common computing component that was routinely left unsecured. It included an observation that was a

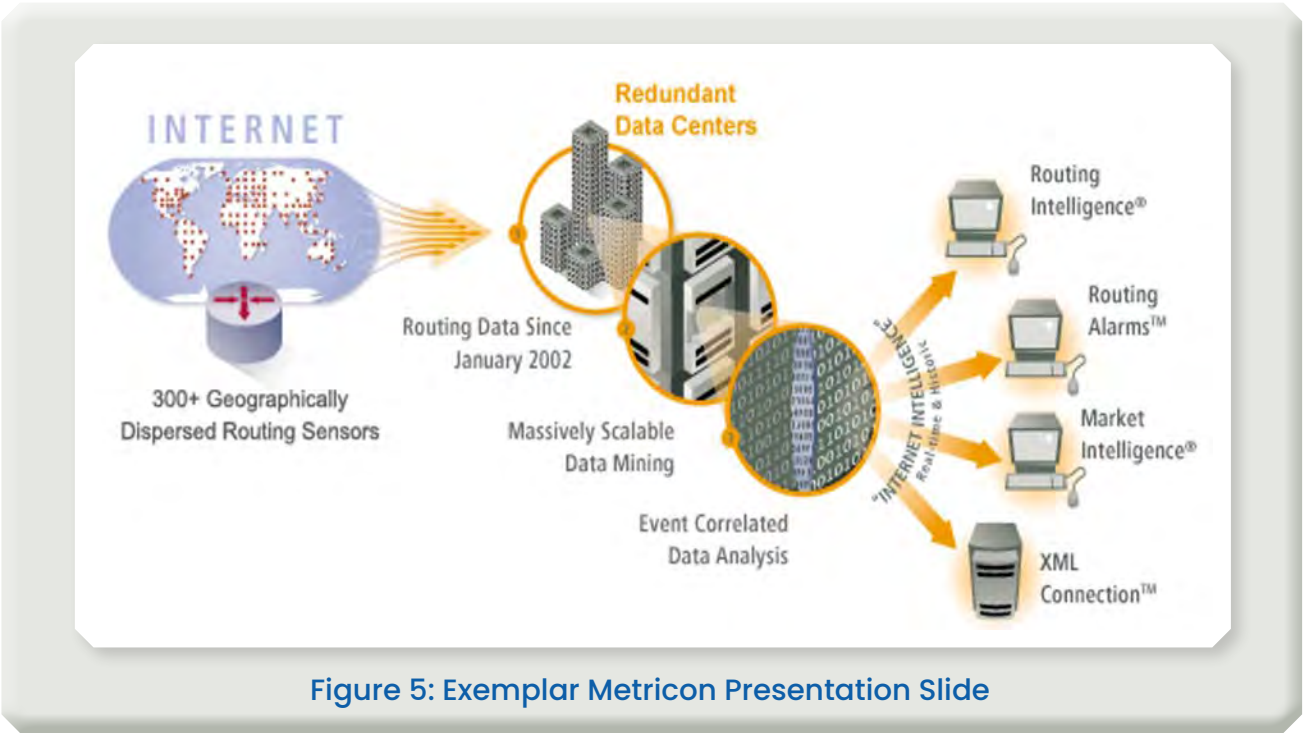

Figure 5: Exemplar Metricon Presentation Slide

recurring theme at Metricon: "Very few people understand these risks, so they are not being measured or managed appropriately. No one is covering your back!" Figure 5 is an excerpt from that presentation. The figure is also an exemplar in that Metricon participants spent a lot of time educating each other on the depth and breadth of technology architecture that must be included in the scope of any enterprise security measurement program.

The many topics discussed at Metricon included: deterministic models, external activity, internal activity, inventory targets, process monitors and threat attributes. As consensus matured, Metricon attendees contributed numerous textbooks, technology articles and academic papers on the topic of measuring cybersecurity.

Metricon researchers advanced the concept of cybersecurity measurement to the point where vendors started selling it. Security scorecard products surfaced in the 2010s. Most of these products scan a company's registered domains from the internet and provide a rating, often numeric, that reflects their view of the company's cybersecurity posture. Scorecard vendors are entirely unregulated, and when they first started, they would rate companies that were not customers as an unsolicited marketing attempt. Unfortunately, vendors did not always identify the target company's domains accurately and sometimes misinterpreted their own scan results. They also typically refused to release their scoring algorithms. By 2017, enough big companies had been given unfairly poor scores to motivate them to mobilize. They brought in the U.S. Chamber of Commerce and came up with security scorecard standards. The ideas was that scorecard vendors that do not follow the **standards** (summarized in Figure 6) could be accused of libel.

The security scorecard marketplace has evolved as corporate management started demanding that cybersecurity be managed in a manner similar to other disciplines.

SECURITY SCORECARD VENDORS SHOULD PROVIDE:

- Transparency into scoring methodology and types of data used;
- Procedures that organizations may use to dispute their score;
- Evidence of accuracy and validation of scoring methodologies and historical performance of scoring models;
- Model governance to include change control over scoring methodologies;
- View and dispute procedures for rated organizations irrespective of whether they are a customer; and
- Confidentiality for scores and for information disclosed by a rated organization during the course of a dispute.

Figure 6:
Summary of Security Scorecard Standards

The security scorecard marketplace has evolved as corporate management started demanding that cybersecurity be managed in a manner similar to other disciplines. As Drucker and Deming pupils would say, "management by observation and control" or "plan-do-study-act." CISOs are routinely required to produce data to support decisions related to a very diverse set of business activities ranging from security operations to board meetings. Details on how to produce metrics are kept in a formal catalog to maintain continuity over staff changes. Increasingly, these metrics-producing activities are classified into categories: intended audience, use cases and production status. Figure 7 is a simple illustration of how CISOs perform such classifications. For anyone engaged in formal cybersecurity metrics programs, this last mention will come as no surprise. We now have metrics on metrics.

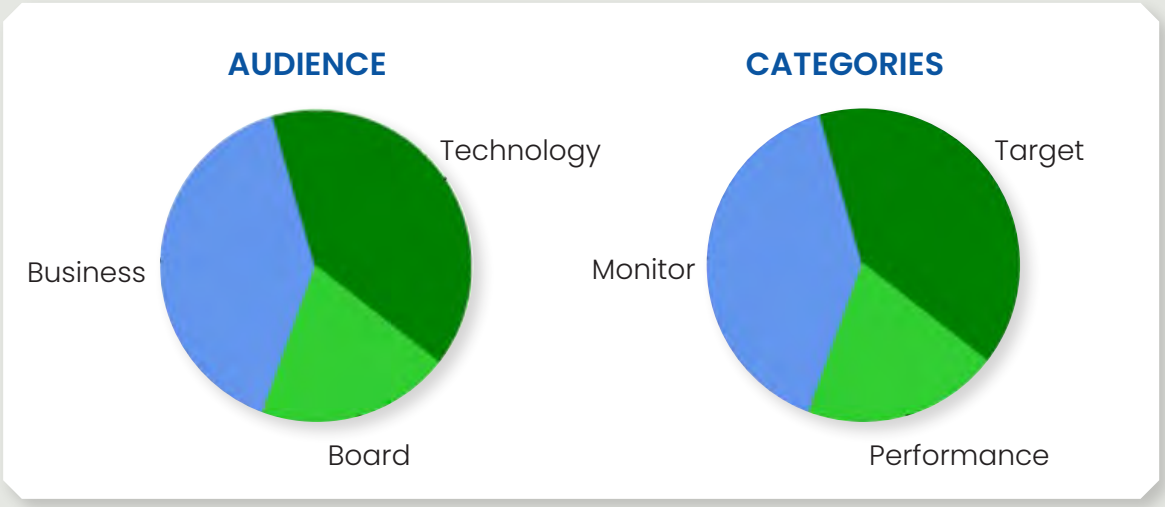

Figure 7: Exemplar Metrics on Metrics

This overview was necessarily brief, so apologies to those who have a favorite cybersecurity metric that was not mentioned. We do not have an industry standard list of security metrics the way the accounting profession has generally accepted accounting principles. Rather, we have business risk appetite guiding customized cybersecurity policies, standards and guidelines. As appropriate in measuring any operational risk, the goal is to find controls, events and issues that support the qualitative risk appetite with quantitative tolerance measures that we can trend over time to see if risk is rising or falling. Just as today's unprecedented fires and floods prompt the weather industry to revisit its assumptions on weather patterns, when metrics indicate risk is falling and we nevertheless succumb to cyberattack, we need to change not only our measures and metrics, but the framework we used to create them.

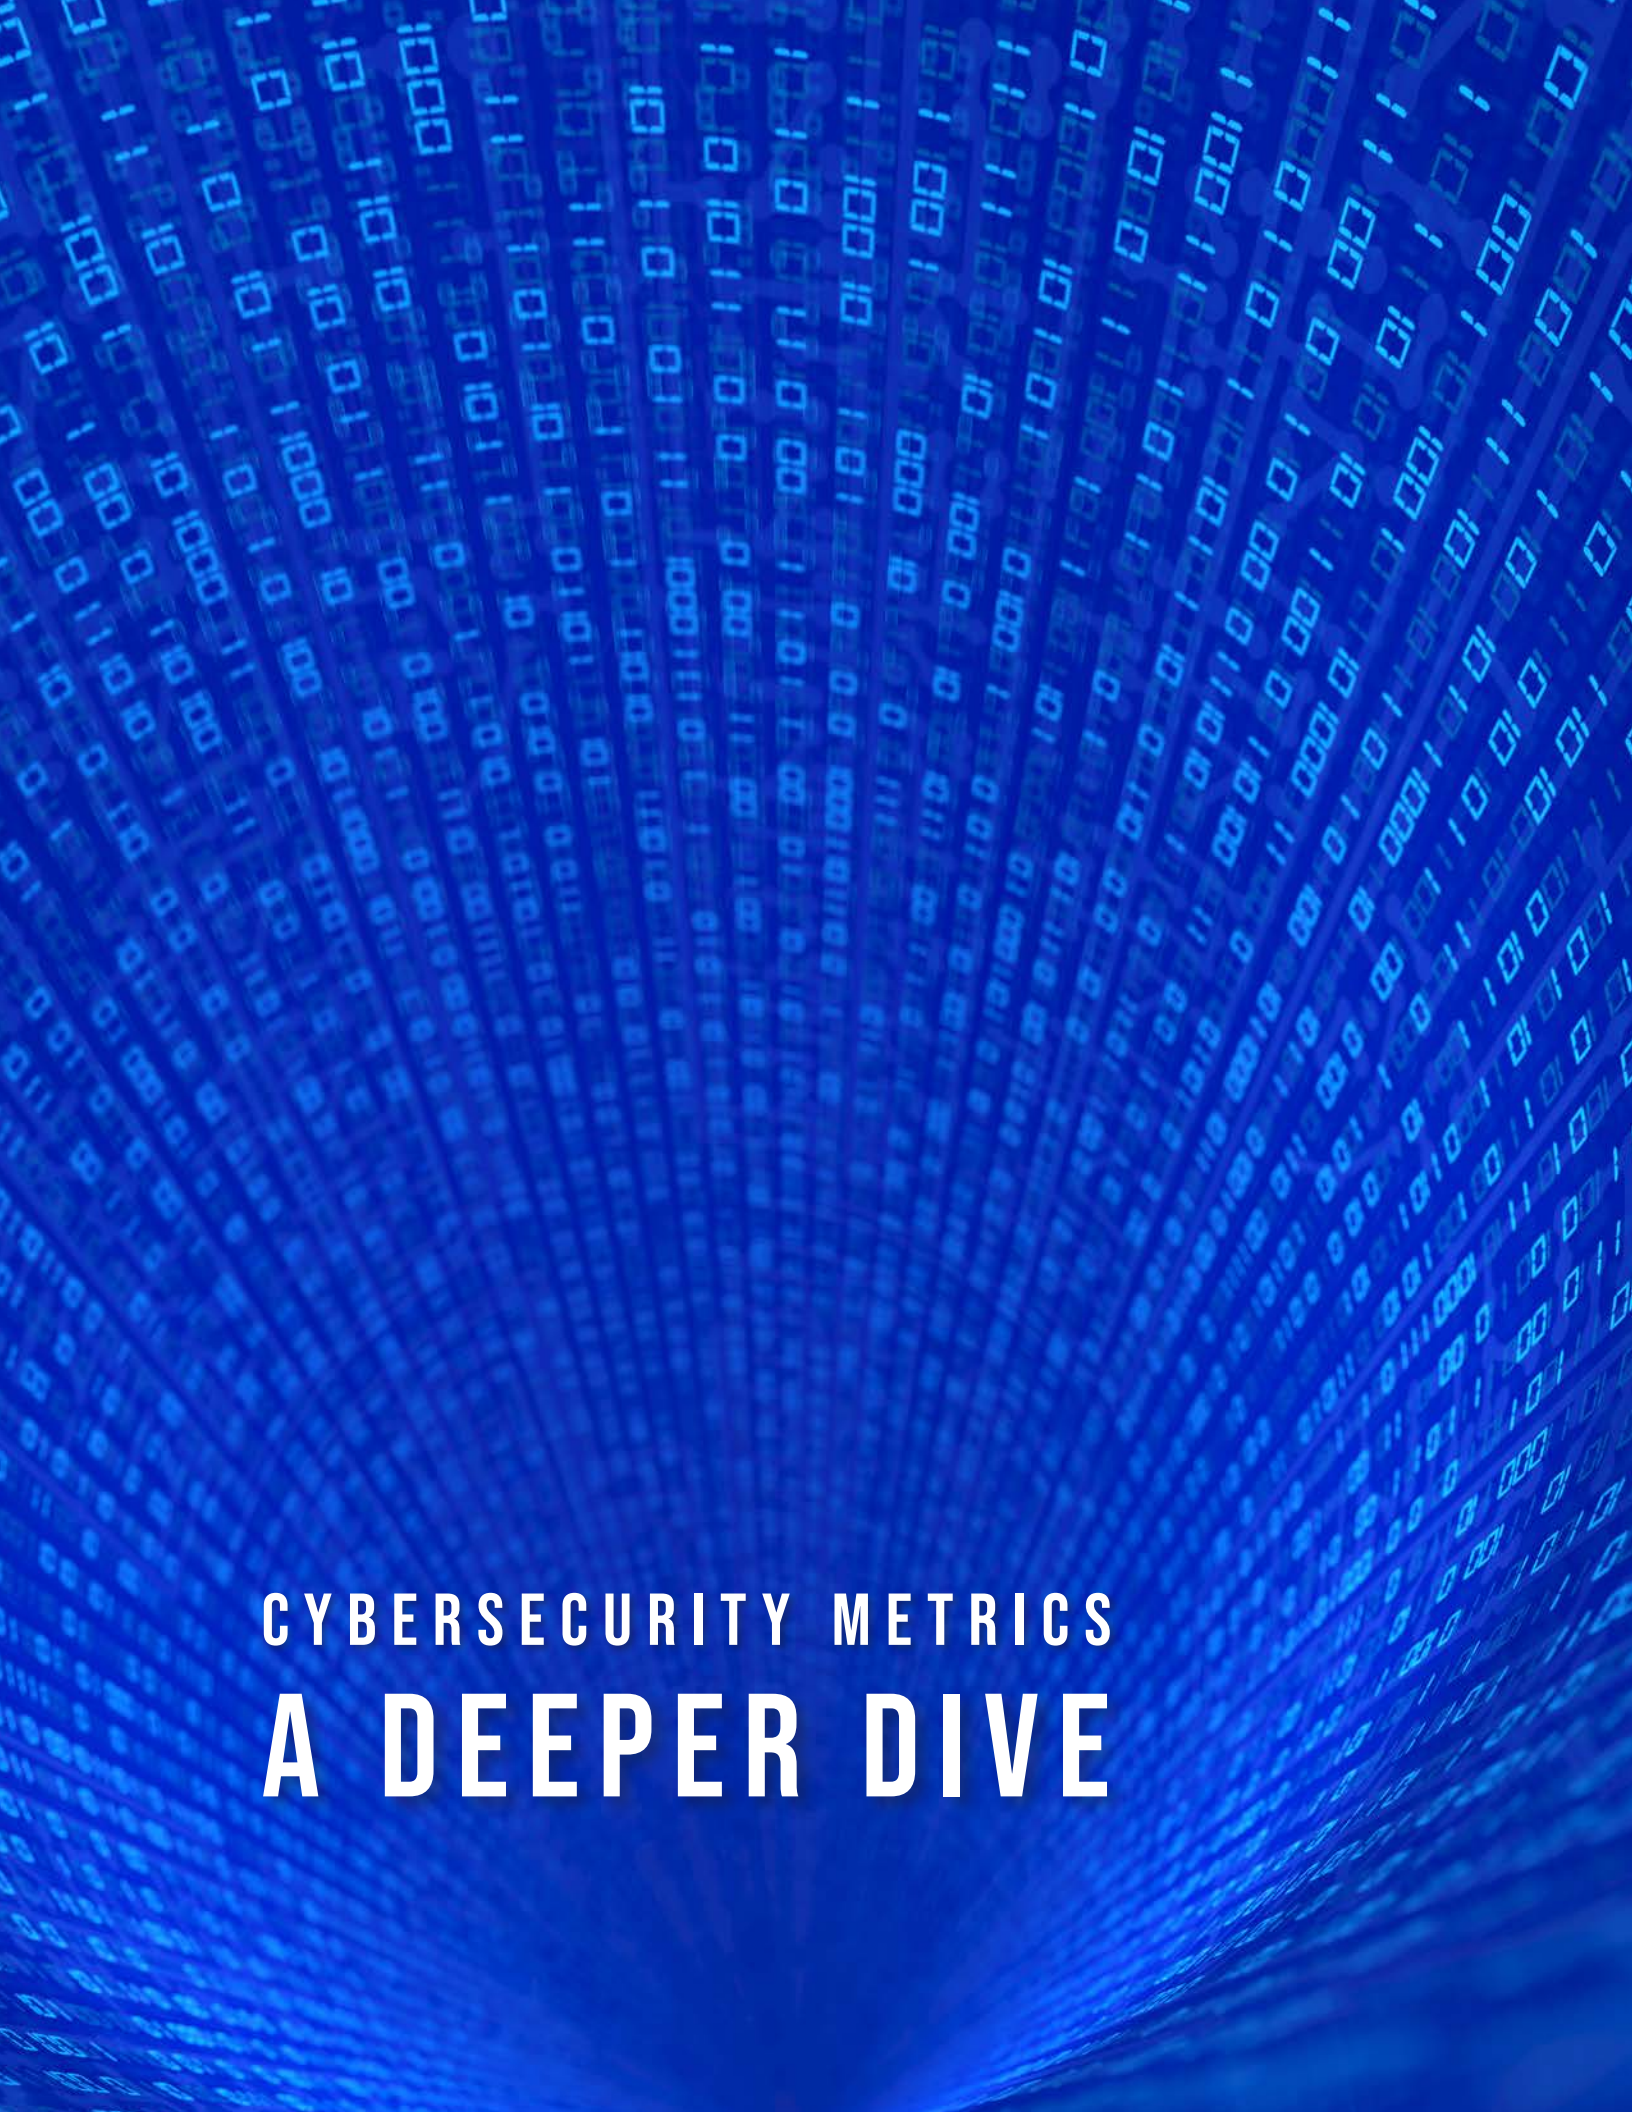

CYBERSECURITY METRICS

A DEEPER DIVE

TAG CYBER

STUDY REPORT

ADVERSARIAL TACTICS, TECHNIQUES & COMMON KNOWLEDGE (ATT&CK) COVERAGE

DR. JENNIFER BAYUK

ATT&CK (Adversarial Tactics, Techniques & Common Knowledge) is a knowledge base of cyberadversary behavior and taxonomy for adversarial actions across their lifecycle. The lifecycle is composed of high-level tactics used by cyberattackers to achieve their objectives, from inception to completion. Each tactic is composed of levels of techniques to complete the phase of the lifecycle. The ATT&CK Framework provides guidance for using ATT&CK to improve enterprise security. Figure 1-1 provides an intuitive feel for how actual cyberattack tactics and techniques are combined to allow enterprises to detect and thwart an adversary's progress through the attack lifecycle. The tactics are the columns and the techniques are in the rows, showing the lifecycle phase identified by the tactic.

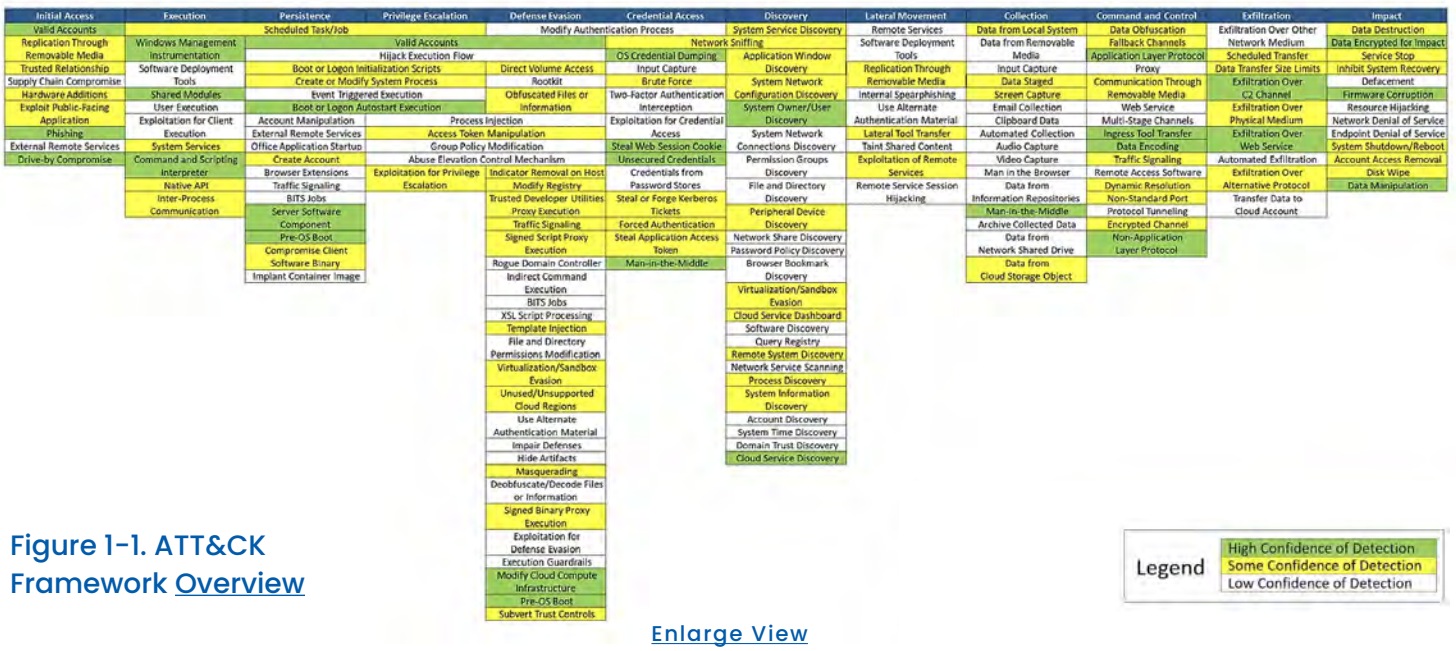

Figure 1-1. ATT&CK Framework Overview

Enlarge View

ATT&CK is used globally by thousands of organizations to measure their cyberdefense preparedness against industry standards. One of those organizations sponsored this study as a means of validating the approach and comparing its performance against cybersecurity industry **practices**. This study has the following objectives:

1.. To provide an independent assessment of the extent to which cyberdefense organizations, engaged in threat detection, have defined as their target the full set of known cyberthreat actor behavior as cataloged by MITRE in its Adversarial Tactics, Techniques & Common Knowledge (ATT&CK®) Knowledge Base.

2.. To provide an analysis of how companies detect cyberattacks so that the results may be used to compare cybersecurity breach defenses across organizations.

This report presents a general overview of the methods by which enterprise cybersecurity teams use ATT&CK to measure threat actor activity within their organizations. We observe that ATT&CK is widely used to both test cyberdefense preparedness and to train security operations on adversary activity. There is no doubt that its use in those efforts is current best practice. We also observe that there is not yet an industry standard—and even less data—on how to measure the extent to which ATT&CK increases the effectiveness of cyberdefense preparedness. Nevertheless, we were able to provide a summary of the current state of these efforts and cite the effectiveness measures that currently make sense. We also include an assessment of ATT&CK's perceived future utility by major corporations.

KEY FINDINGS

The use of ATT&CK as a framework that drives cybersecurity defense requirements has become increasingly common since its introduction in 2013.

- Although the purpose of ATT&CK is to help enterprises measure cyberdefense preparedness, both creators and users of the Framework caution that its measures are not exact or quantitative but inherently qualitative since its methodology requires a series of judgments on whether attack techniques are adequately thwarted.
- Because the ATT&CK data constantly evolves, it is difficult to make comparisons across historical testing. This dearth made it challenging to find data similar enough to meet our objective to compare cybersecurity breach defenses across organizations. Nevertheless, where participants kept track of the successful defense scores over successive exercises utilizing the same attack patterns, all observations were that cyberdefense performance increased significantly (~25-40%). Participants agree that this type of data is a best practice assurance indicator that the exercise results are in fact used to strengthen cyberdefenses.
- Participants agree that the further an attack tactic penetrates into the infrastructure (for example, activity after an initial periphery breach and a lateral move to a different part of the infrastructure), both detection and prevention capability fall considerably. This indicates that later stages of the ATT&CK Framework are not effectively utilized.
- Some study participants observed that our research question was similar to those posed by the **Center for Threat Informed Defense** (CTID). CTID's steady growth since its inception in 2019 reflects the cybersecurity community's embrace of the MITRE Framework.

BACKGROUND

The ATT&CK knowledge base is part of a larger cybersecurity data set, mostly hosted by MITRE within its charter as a U.S. federally funded research and development center for focus on public-private partnership to share best practices on cybersecurity. The larger data set includes:

- **Common Vulnerability Enumeration (CVE):** a collection of known vulnerabilities in technology products.
- **Common Weakness Enumeration (CWE):** a set of known weaknesses in cybersecurity controls.
- **Common Vulnerability Scoring System (CVSS):** a method of risk-rating vulnerabilities based on a given system of interest.
- **Common Attack Pattern Enumerations and Classifications (CAPEC):** links common attack patterns to CWEs.

ATT&CK provides information on which threat actors utilize which attack patterns and maps those patterns to attack vectors. An additional database, the Known Exploited Vulnerabilities Catalog (KEV), hosted by the U.S. Department of Homeland Security's Cybersecurity and Infrastructure Security Agency (CISA), provides notification of current active exploits of CVEs by specific threat actors. This information allows enterprises to prioritize cyberdefense efforts based on the corresponding increase in cyberattack risk.

Cybersecurity breach and attack platforms typically incorporate all of this data into their products, linking the data sets and allowing chief information security officers (CISOs) to prioritize attack simulation and training based on their own vulnerability in combination with expected attacks. ATT&CK provides the basis for realistic attack simulation and enables detection tools to recognize that individual suspicious activities detected across an enterprise network are actually indicators of a single attack path. The flow in which this data is accumulated and delivered is depicted in Figure 3-1.

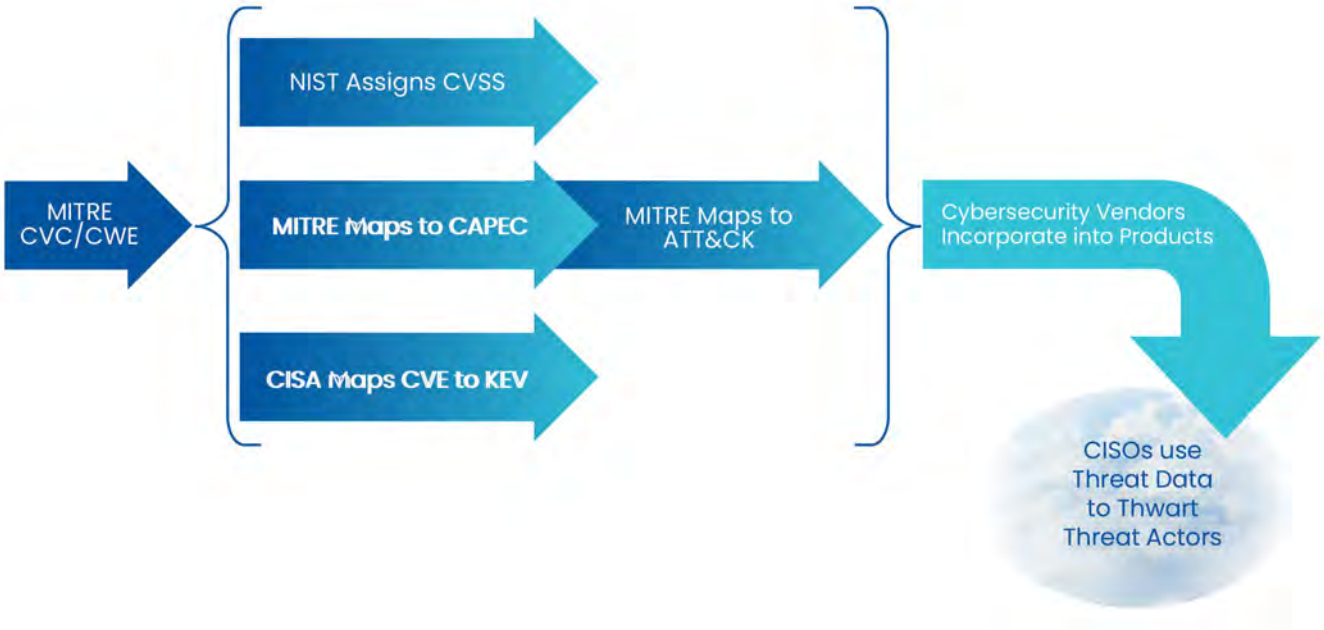

Figure 3-1. ATT&CK Use Case Data Flow

Note that Figure 3-1 does not contain the workflow used by a CISO to incorporate the available data sets into their cyberdefense measurements. A CISO must first envision their own enterprise infrastructure

business application landscape before it makes sense to select puzzle pieces from ATT&CK to help fill in visibility gaps. In the midst of this activity, new versions of ATT&CK may increase the number of attack path tactics, new threat actor techniques are documented and added to those associated with a given tactic, and there is a time gap between discovery of new threat actor techniques and ATT&CK update. The vendors know this, and have all created their own classification systems that incorporate MITRE ATT&CK rather than try to use it literally. In a training course intended for the assessment of cyberdefense capability using ATT&CK, MITRE itself cautions that such assessment "paints broad strokes to give you an indication of where—generally—**your gaps lie**" and that "Coverage inference up/down the hierarchy is almost always dependent on the context and user preference." Note that we did not evaluate coverage accuracy, we simply consolidated and reported data provided by participants. Many vendors cited the fuzziness of not just ATT&CK, but the links between the data upon which it is based. One participant shared a study on the topic that summarized the situation (Figure 3-2).

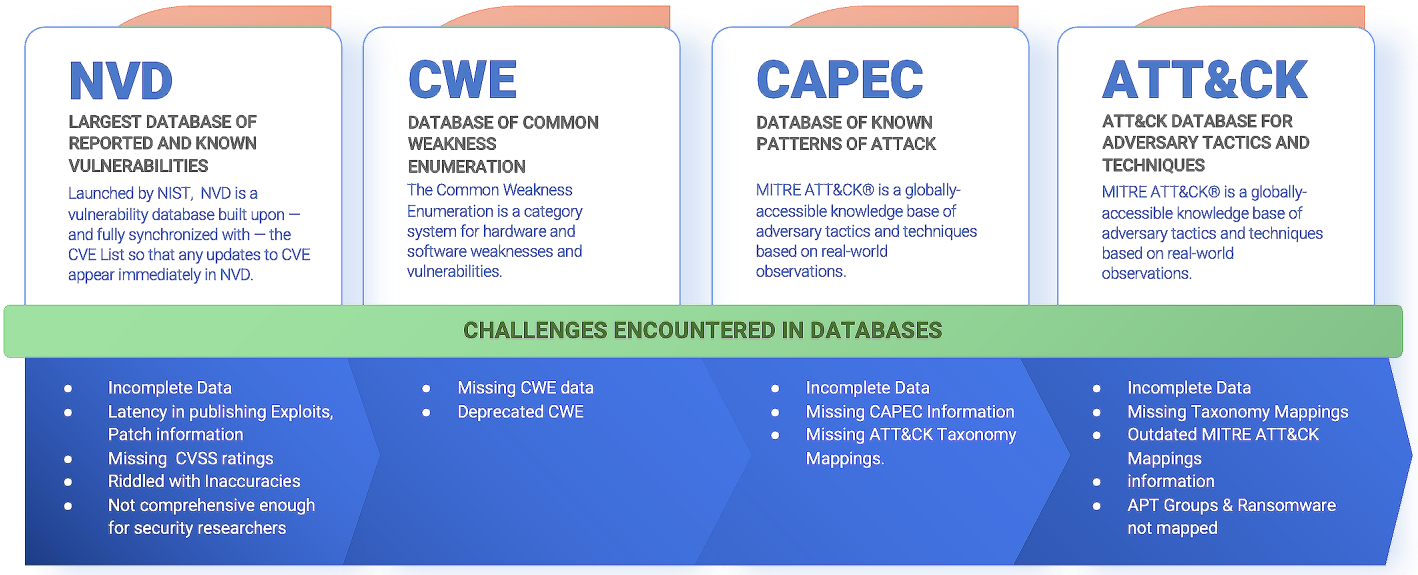

Figure 3-2. Challenges to Precise Enterprise Mapping to ATT&CK

So although there is a common methodology across firms that utilize MITRE, there is no common method, which made it difficult to compare results across organizations that use the ATT&CK Framework. Even if it were not constantly changing, no organizations are on the same implementation timeframe. Therefore, the scope of the Framework used for this study is the least common denominator among its participants (see Appendix C for details on scope).

Nevertheless, the situation does not appear to dissuade the cybersecurity community from rallying around ATT&CK as the best source of hope for future preparedness. Many vendors cited this hope as their reason for participating in this study. Some participants had already joined forces with MITRE by sponsoring various extensions to ATT&CK development to close the gaps. CTID project participation further confirms community commitment to develop ATT&CK over inventing alternatives. CTID projects include collaboration to map adversary activity, developing ATT&CK matrices for new technologies, mapping ATT&CK to control standards and detection data sources, and developing tools for analyzing ATT&CK data for use by security operations.

RESEARCH METHOD

The research method follows the standard strategy used by most senior research analysts: begin with a research question, develop a hypothesis, gather and analyze relevant measures to test the hypothesis, and form a conclusion.

The research question was provided by the study's sponsor:
To what extent are the tools and techniques in active use by cyberdefense organizations in major corporations capable of detecting and mitigating the threats outlined in the ATT&CK Library?

The corresponding research hypothesis is:
Cyberdefense organizations in major corporations are capable of detecting and mitigating the threats outlined in the MITRE ATT&CK Library, if and only if their technology platforms are sufficiently instrumented to gather and analyze the data require to detect exploits identified in the ATT&CK Library.

The measures gathered to test this hypothesis were sourced from cybersecurity vendors that provide technology platforms instrumented to gather and analyze the data required to detect exploits identified in the ATT&CK Library. These are primarily breach and attack simulation (BAS) tools. Enterprise cyberdefense teams utilize BAS tools to provide continuous validation of certain deployed cybersecurity controls. Therefore, the TAG Cyber solicitation focused on all vendors classified by TAG Cyber as belonging in the BAS category of the TAG Cybersecurity Taxonomy. Discussions were held with participants to answer their questions on the study and to learn their approach to integrating ATT&CK into their products and services. Participants then sent TAG Cyber samples of data and/or reports that they thought were germane to the study's objectives. TAG Cyber analyzed and consolidated this data to produce this report.

PROJECT EXECUTION

An invitation to participate was sent via email to existing TAG vendor contacts; if none existed, the email was sent to marketing leads at the firm. In all, 19 vendors were invited. Most of the potential participants were BAS vendors. A few were not strictly BAS, but incorporated the MITRE Framework into their products for MITRE coverage mapping. For example, where TAG found that vendors in related cybersecurity fields had previously published statistics on MITRE coverage, these were included.

Eleven (11) of the potential participants responded positively to the initial invitation and shared their thoughts via a discussion and/or data submission. Most of the data was in the form of previously published reports or presentations that contained their own analysis of MITRE coverage and other observations on use cases for the MITRE ATT&CK Framework. All data was anonymized except for the participant data source. Most were based on interactions with their customer product use cases, but some were based on surveys or other non-product-related customer interaction. These vendors are listed in Appendix B.

Two weeks before the deadline for our first draft, reminders were sent to those that did not yet submit data. Two reminders were sent to those that initially responded positively but had not submitted data.

The use cases for ATT&CK mapping are widespread, ranging from threat intelligence to third party visibility.[1] We focused on the use case of strengthening the capability of a security operations center to impede threats. This is typically done by using ATT&CK to identify the expected activities of threat actors that may target the firm. This creates requirements for that activity to be visible to the security operations center and mapped onto an attack tactic (attack lifecycle phase). In combination with

[1] *For example, these two reports include completely different top-ten use cases for ATT&CK lists:* www.attackiq.com/lp/mitre-attack-for-dummies/ *and* www.splunk.com/en_us/form/10-ways-to-take-the-mitre-att-and-ck-framework-from-plan-to-action.html

vulnerability monitoring of the targeted technology, the attack may be visibly thwarted or successful. The number and type of tactics and techniques that are tested, in combination with test results, provide the basis for a measurable comparison (see Figure 4.1-1).

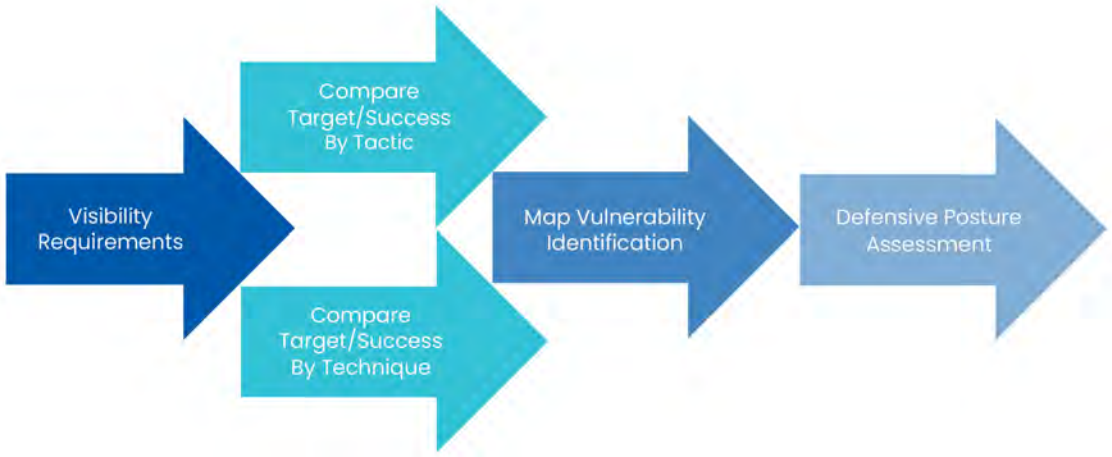

Figure 4.1-1. Use Case ATT&CK Data Flow

ATT&CK COVERAGE COMPARISON

We appreciated all the materials and insight shared by the study participants. We found that nine (9) of the ten (10) participants used ATT&CK internally to catalog the attack patterns with their software and the 10th will be introducing that feature within the next few months. There was enough similarity across seven (7) of the participants to compare the extent to which ATT&CK tactics were included in their statistics and observations. However, due to the inherent issues in aggregating data gathered from independent processes with different objectives, as well as confidentiality concerns, only four (4) participants shared enough of their measures to directly assess targeted ATT&CK coverage.

At the start of the study, every participant was willing and eager to share the results of their own analysis, but for various reasons few could get permission to share their actual data sets. Nevertheless, the comparisons in our results are significant in that:

- We confirmed that methods of testing defense preparedness using MITRE ATT&CK is cybersecurity industry best practice;
- The use of ATT&CK as a standard across all participants enabled us to make good use of highly relevant subset data, and the increasing availability of these studies ensures that notable subsets should inform security program metrics going forward;
- Tracking successful defense scores over successive exercises utilizing the same attack patterns is a best practice assurance indicator that the exercise results are in fact used to strengthen cyberdefenses; and
- Participants were introduced to a peer group that endorses ATT&CK data sharing so will have more opportunities for peer measurement going forward.

TARGET COVERAGE

While none of the organizations included all the techniques within our scope (as defined in Appendix C), there was enough similarity across our four data-sharing participants for us to compare the extent to which their testing of the techniques within our scope was successful. Where more detailed data was available at the tactics and/or techniques level, lower level data was aggregated to a allow comparison at higher levels. The result is a set of percentages that reflects the extent to which attacks were thwarted in simulated environments (see Figure 5.1-1).

The underlying measures based on the ATT&CK technique level are:

A. Number of distinct tests of an ATT&CK technique for which the target blocks successful execution of the technique.

B. Number of distinct tests in A that were successfully blocked.

C. Number of distinct tests of an ATT&CK technique for which the target detects execution of the technique.

D. Number of distinct tests in C that were successfully detected.

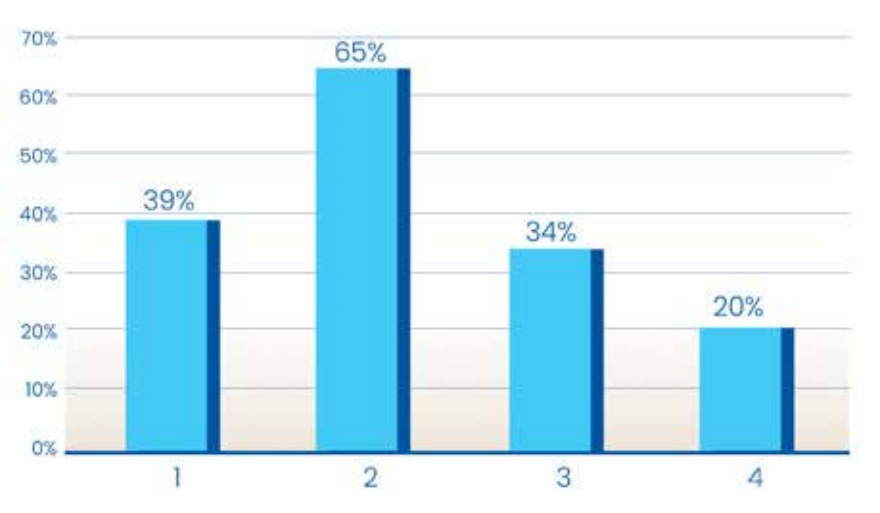

Figure 5.1-1. ATT&CK Coverage

That is, each technique was assigned a measure calculated as (A+C)/(B+D).

The underlying measures based on the ATT&CK tactic level were then calculated to gain the average percentage of passed technique tests for each tactic in scope. Figure 5.1-1 shows the customized target as 100%. The score displayed is the extent to which the target is met.

Time periods for the tests varied but were based on a 12-month history. The percentages thus correspond to enterprise-relative success in detecting and preventing attacks included in the set of ATT&CK techniques for which organizations want to be prepared (albeit "in broad strokes" as per our discussion in section 3). Several participants acknowledged that enterprises increase capability for successful defense after a few simulations of the same ATT&CK techniques, so current performance may be higher than the averages displayed. However, it is also the case that, as new systems and organizations join the simulation pool, the average can dip significantly. Nevertheless, the data is relevant because (i) technology itself is constantly changing and (ii) only the most cybersecurity-mature companies even attempt such simulation. It is expensive and time-consuming, making it out of reach to most medium and small businesses with more stable technology environments.

MOST FREQUENT ATTACK COVERAGE

One of our participants recently conducted an analysis of the extent to which organizations were able to cover the most frequent attacks seen in the wild. We took advantage of their ATT&CK map to compare the participant's coverage of the most frequent attacks using the technique measure described in Section 5.1. The result is a much closer distribution of success levels, indicating that those most frequent attack techniques were likely included in the targets of the peer comparison participants. Figure 5.2-1 displays that result.

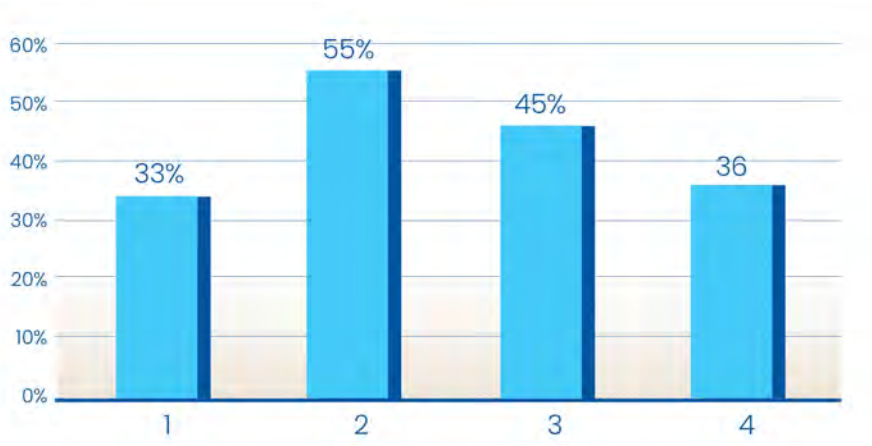

Figure 5.1-2. Most Frequent ATT&CK Coverage

TOP TEN EXPLOITS

One of our participants is a post-attack forensics service provider. This provider does not collect ATT&CK testing data, but had a relevant contribution because it recently conducted an analysis of trends in attacks under investigation. We took advantage of this ATT&CK map to compare other participants' coverage of the top ten exploits listed. The result, depicted in Figure 5.3-1, moved away from the convergence seen between Figures 5.1-1 and 5.1-2. Instead, the differences between Figures 5.1-2 and 5.1-3 are more pronounced. Were this result to be replicated in a larger study, it would indicate that the actually exploited top threats may not be routinely included in the least common denominator of BAS tests, and the more successful outliers may have more reliable predictions of successful adversary behavior.

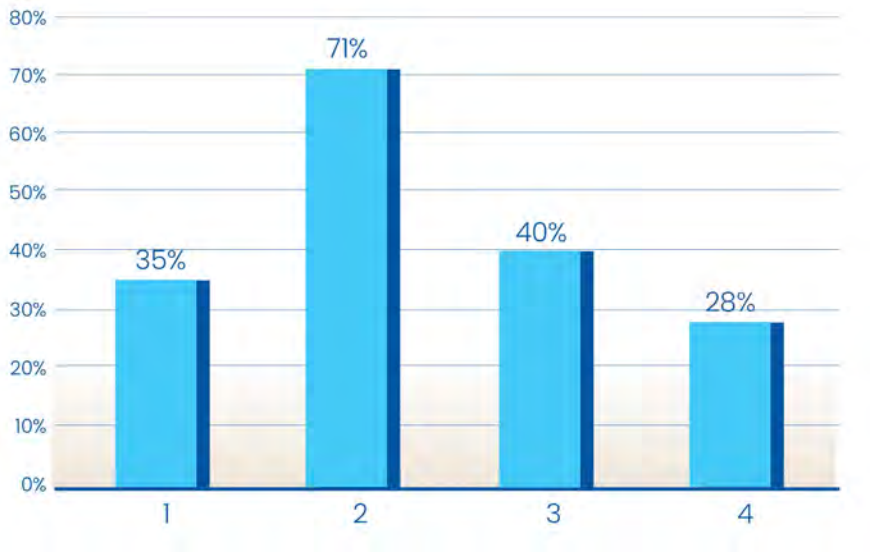

Figure 5.1-3. Top Ten Exploits ATT&CK Coverage

CONCLUSION

The study results showed that our research hypothesis is true. The cyberdefense organizations in major corporations have technology platforms that are sufficiently instrumented to gather and analyze the data to detect exploits identified in the ATT&CK Library. Hence, they are capable of detecting and mitigating the threats outlined in the ATT&CK Library. Therefore, we were able to use this data to approximate an answer to the research question on the extent to which these tools and techniques are capable of detecting and mitigating the threats outlined in the ATT&CK Library; roughly between 20 and 65%.

Of course, any tool is only as effective as its operator, and our study is hampered by a limited data set. Hence, we fall back on a traditional disclaimer: If there is additional evidence available to be examined at a later date, our opinion may change. It is therefore our recommendation that those interested in more exact answers to these questions join peer study groups that are already focused on the issue (such as CTID), or use their existing information-sharing communities to start an initiative based on ATT&CK patterns (e.g. via **industry Information Sharing and Analysis Centers**).

The ATT&CK Framework may not be perfect when it comes to measuring cyberdefense preparedness, but clearly it is the best current option. The ten vendors we interacted with during the study had different perspectives on the SIEM and BAS product architecture, features and functions, yet they independently arrived at the conclusion that ATT&CK was a necessary component of their data model going forward. That we were able to take advantage of their research and interpret others' data based on it is evidence of the Framework's utility.

Although industry standards in ATT&CK application are not prescriptive, the community is converging on an assurance-level approach to defense preparedness using ATT&CK. This circumstance is common in the cybersecurity profession, which is still very young and yet to converge into norms of operation. We expect ATT&CK use cases to evolve with increased collaboration and automation fueled by public-private partnerships with the common goal of national defense.

APPENDIX A: EXAMPLE INVITATION TO PARTICIPATE

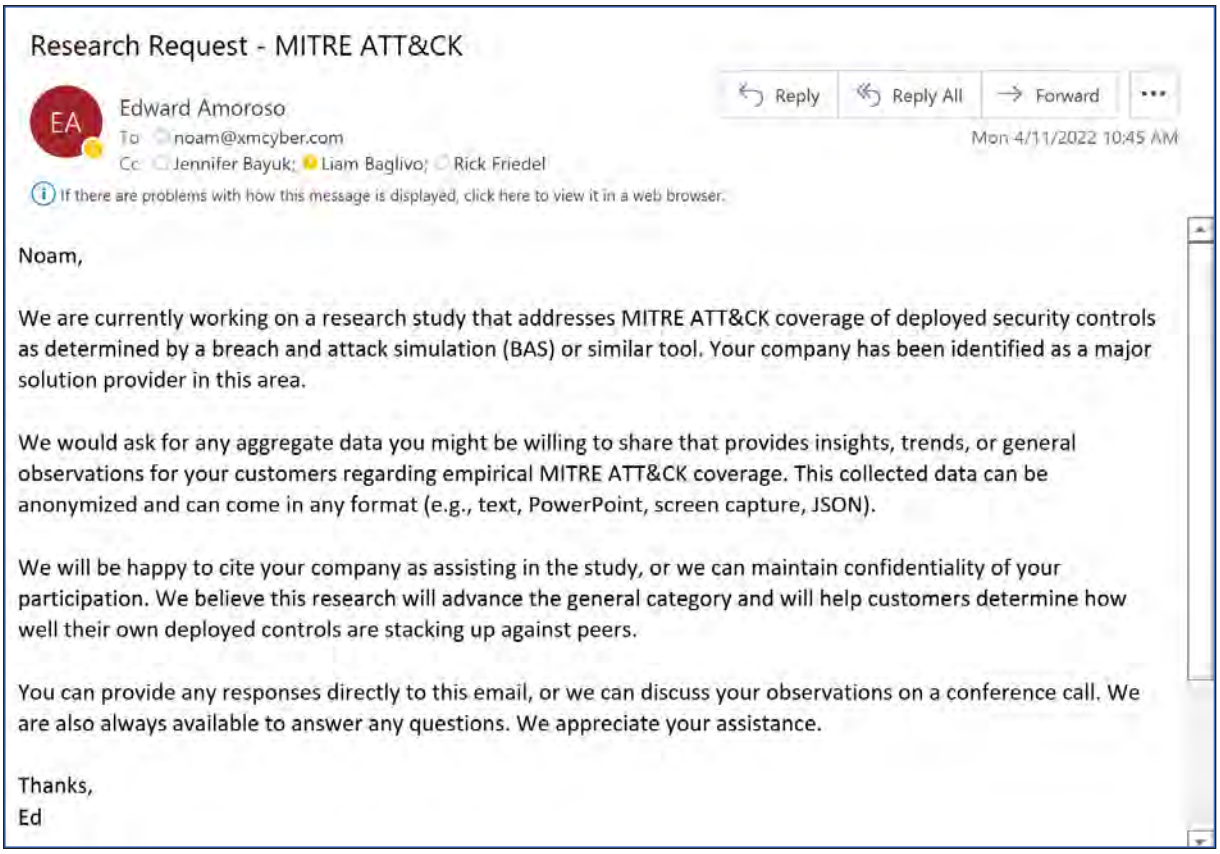

Research Request - MITRE ATT&CK

Reply | Reply All | Forward | ...

EA

Edward Amoroso

To noam@xmcyber.com

Cc Jennifer Bayuk; Liam Baglivo; Rick Friedel

Mon 4/11/2022 10:45 AM

If there are problems with how this message is displayed, click here to view it in a web browser.

Noam,

We are currently working on a research study that addresses MITRE ATT&CK coverage of deployed security controls as determined by a breach and attack simulation (BAS) or similar tool. Your company has been identified as a major solution provider in this area.

We would ask for any aggregate data you might be willing to share that provides insights, trends, or general observations for your customers regarding empirical MITRE ATT&CK coverage. This collected data can be anonymized and can come in any format (e.g., text, PowerPoint, screen capture, JSON).

We will be happy to cite your company as assisting in the study, or we can maintain confidentiality of your participation. We believe this research will advance the general category and will help customers determine how well their own deployed controls are stacking up against peers.

You can provide any responses directly to this email, or we can discuss your observations on a conference call. We are also always available to answer any questions. We appreciate your assistance.

Thanks,
Ed

APPENDIX B: STUDY PARTICIPANTS

This study was sponsored by UnitedHealth Group, a multinational corporation headquartered in Minnesota which focuses on managed healthcare and insurance: www.unitedhealthgroup.com

The study was conducted by TAG Cyber LLC, a global research and advisory firm headquartered in New York that focuses on cyber security issues for enterprise, government, and commercial vendors: www.tag-cyber.com

Both UnitedHealth Group and TAG Cyber sincerely appreciate the advice, guidance and/or data contributions from the following cybersecurity vendor community study participants:

- AttackIQ
- CardinalOps
- Cloud Range
- Cymulate
- Ivanti
- Mandiant
- Picus
- SafeBreach
- SCYTHE
- XM Cyber

APPENDIX C: STUDY SCOPE

Our analysis focuses on ATT&CK for Enterprise. There are also versions for Mobile and Industrial Control Systems. Versions for Cloud and Container are planned.

Our scope includes Tactics 1-11 and Tactic 40. The other two Tactics, 42 and 43, have been more recently introduced an include techniques used by actors prior to committing an actual attack on an enterprise.

ATT&CK Tactics – Listed in order of Appearance in Attack Path			
Introduced	Number	Name	Description
Oct 2020	TA0043	Reconnaissance	The adversary is trying to gather information they can use to plan future operations.
Oct 2020	TA0042	Resource Development	The adversary is trying to establish resources they can use to support operations.
Oct 2018	TA0001	Initial Access	The adversary is trying to get into your network.
Oct 2018	TA0002	Execution	The adversary is trying to run malicious code.
Oct 2018	TA0003	Persistence	The adversary is trying to maintain their foothold.
Oct 2018	TA0004	Privilege Escalation	The adversary is trying to gain higher-level permissions.
Oct 2018	TA0005	Defense Evasion	The adversary is trying to avoid being detected.
Oct 2018	TA0006	Credential Access	The adversary is trying to steal account names and passwords.
Oct 2018	TA0007	Discovery	The adversary is trying to figure out your environment.
Oct 2018	TA0008	Lateral Movement	The adversary is trying to move through your environment.
Oct 2018	TA0009	Collection	The adversary is trying to gather data of interest to their goal.
Oct 2018	TA0011	Command and Control	The adversary is trying to communicate with compromised systems to control them.
Oct 2018	TA0010	Exfiltration	The adversary is trying to steal data.
Mar 2019	TA0040	Impact	The adversary is trying to manipulate, interrupt, or destroy your systems and data.

In Scope Number of Techniques 326
Number of Sub-Techniques 379
Number of Techniques mapped to Enterprise 233
Number of Enterprise Techniques in Scope 216
Number of Enterprise Sub-techniques in Scope 454

For the Most Frequent Attack Coverage, these 14 Techniques were provided by CardinalOps:

T1059 Command and Scripting Interpreter
T1218 Signed Binary Proxy Execution
T1543 Create and Modify System Process
T1053 Scheduled Task / Job
T1027 Obfuscated Files or Information
T1105 Ingress Tool Transfer
T1569 System Services
T1036 Masquerading
T1486 Data Encrypted for Impact
T1082 System Information Discovery
T1497 Virtualization/Sandbox Evasion
T1098 Account Manipulation
T1219 Remote Access Tools
T1018 Remote System Discovery

For the Top Ten Forensics Investigations, these 11 Techniques were provided by Mandiant (there are 11 because there was a tie for #10):

T1027 Obfuscated Files or Information
T1059 Command and Scripting Interpreter
T1071 Application Layer Protocol
T1082 System Information Discovery
T1083 File and Directory Discovery
T1070 Indicator Removal on Host
T1055 Process Injection
T1021 Remote Services
T1497 Virtualization/Sandbox Evasion
T1105 Ingress Tool Transfer
T1569 System Services

APPENDIX D: REFERENCES

ATT&CK Background

Adversarial Tactics, Techniques & Common Knowledge (ATT&CK), National Cybersecurity FFRDC (MITRE),

Center for Threat-Informed Defense (CTID),

Common Attack Pattern Enumerations and Classifications (CAPEC), U.S. Department of Homeland Security (MITRE),

Common Vulnerability Collaboration CVE Numbering Authorities,

Common Vulnerability Enumerations, National Cybersecurity FFRDC (MITRE),

Common Weakness Enumeration, National Cybersecurity FFRDC (MITRE),

Living Off The Land Binaries, Scripts and Libraries (LOLBAS), Open source collaboration to collect and use ATT&CK to classify binaries that can be used by an attacker to perform actions beyond their original purpose.

U.S. Cybersecurity and Infrastructure Security Agency, Known Exploited Vulnerabilities Catalog,

Vulnerability Map and Common Vulnerability Scoring Systems (CVSS), U.S. National Institute of Standards and **Technology National Vulnerability Database (NVD)**

Participant Contributions

AttackIQ	**Center for Threat-Informed Defense 2022 Impact Report**
CardinalOps	**2022 Report On State of SIEM Detection Risk**
Cymulate	**2021 State of Cybersecurity Effectiveness Report**
Ivanti	**2021 Zero Trust Progress Report**
Mandiant	**2022 Mandiant Trends Report**
Picus	**The Red Report 2021**
SafeBreach	**Operationalize MITRE ATT&CK™ Framework,**
SCYTHE	**The Purple Team Framework**
XM Cyber	**Cyber 2022 Attack Path Management Impact Report**

A SENTIMENT-BASED INDEX TO MEASURE THE CYBERSECURITY THREAT

ANDY MCCOOL

What if you could get a clear picture of the broad cybersecurity landscape? What if you could get real world insight into the threats, attacks and weapons being used? What if you could understand the trends being seen by others in the industry? It is with these goals in mind that we have been working with the New York University Index of Cybersecurity (ICS) to see if we can develop cybersecurity metrics that can be used to provide these insights.

BRIEF HISTORY

The ICS is a sentiment-based index that was created in 2011 by Dan Geer and Mukul Pareek. At the time, Geer was (and still is) the CISO of In-Q-Tel, which invests in cutting-edge technologies. He referred to himself as a "security researcher with a quantitative bent." Pareek, who is now a senior vice president at Wells Fargo and focuses on technology control modeling and analytics, called himself "a risk management professional." They described the project this way: "Monthly, we poll people with operational responsibility for cybersecurity on how two dozen different cybersecurity risks have changed in the past month, and, from that polling, calculate the ICS." The results show whether cybersecurity threats have increased or declined.

In 2018, the management and operation of the ICS was transferred to a research team led by Dr. Edward Amoroso at New York University's Tandon School of Engineering. He is assisted by experts here at the advisory firm TAG Cyber (where Amoroso is the CEO and founder and where I also work). The move to NYU was designed to reinvigorate involvement by the contributors, to analyze the data in search of useful correlations and to investigate the possibility of creating an improved series of parallel ICS-related metrics moving forward.

INTENT AND METHODOLOGY

The ICS provides a numeric estimate of cybersecurity posture based on the collective sentiment of the select practitioners using a monthly email survey. The survey consists of six questions that cover attack actors, weapons, effects desired by attackers, attack targets, defenses and overall perception. Each question has subquestions to measure a specific component of the topic

area. An additional question of the month is included to solicit views on a current cybersecurity topic. The high-level results are posted each month to the ICS website (https://cybersecurityindex.org) and a more detailed analysis is sent to the participants.

The selection criteria used to choose expert responders recognizes that survey research can be susceptible to misleading information from less experienced participants. To address this, the responder list comprises industry practitioners with known competence, experience and expertise. They include chief information security officers (CISOs), direct reports to CISOs, academics, technology product vendors' technical experts or chief scientists, and chief information risk officers. Responses from experts are collected without attribution. That is, their responses are sent to the ICS administrative team, but information about their email addresses is neither logged nor analyzed. In fact, metadata about expert responders (such as industry, years of experience and management level) is not retained.

INDEX CALCULATION

The absolute cumulative aggregate value of the threat index reflects not only the sentiment offered on a given month, but also factors in all previous reporting periods. This is accomplished through continuous compounding of the index value—with the goal that an increase by some percentage in one month followed by a decrease of the same percentage would return the absolute index value to the original value. The general purpose is to track whether the combined sentiment of responders is going up, going down, or staying the same on a month-to-month basis.

The Index uses a five-level Likert scale for scoring answers, with all subquestions in any given area weighed equally to produce the aggregate subtotal score. The percentages assigned to the five-level answers are as follows: -20% (fallen fast), -7.5% (fallen), 0% (stayed static), +7.5% (risen) and +20% (risen fast). The subtotal scores for each of the subindex questions are added together and divided by the number of questions to obtain the score from a given responder. This allows averaging all responses for a given month and using the result for the index calculation. The ICS results for August 2022 are depicted in Figure 1.

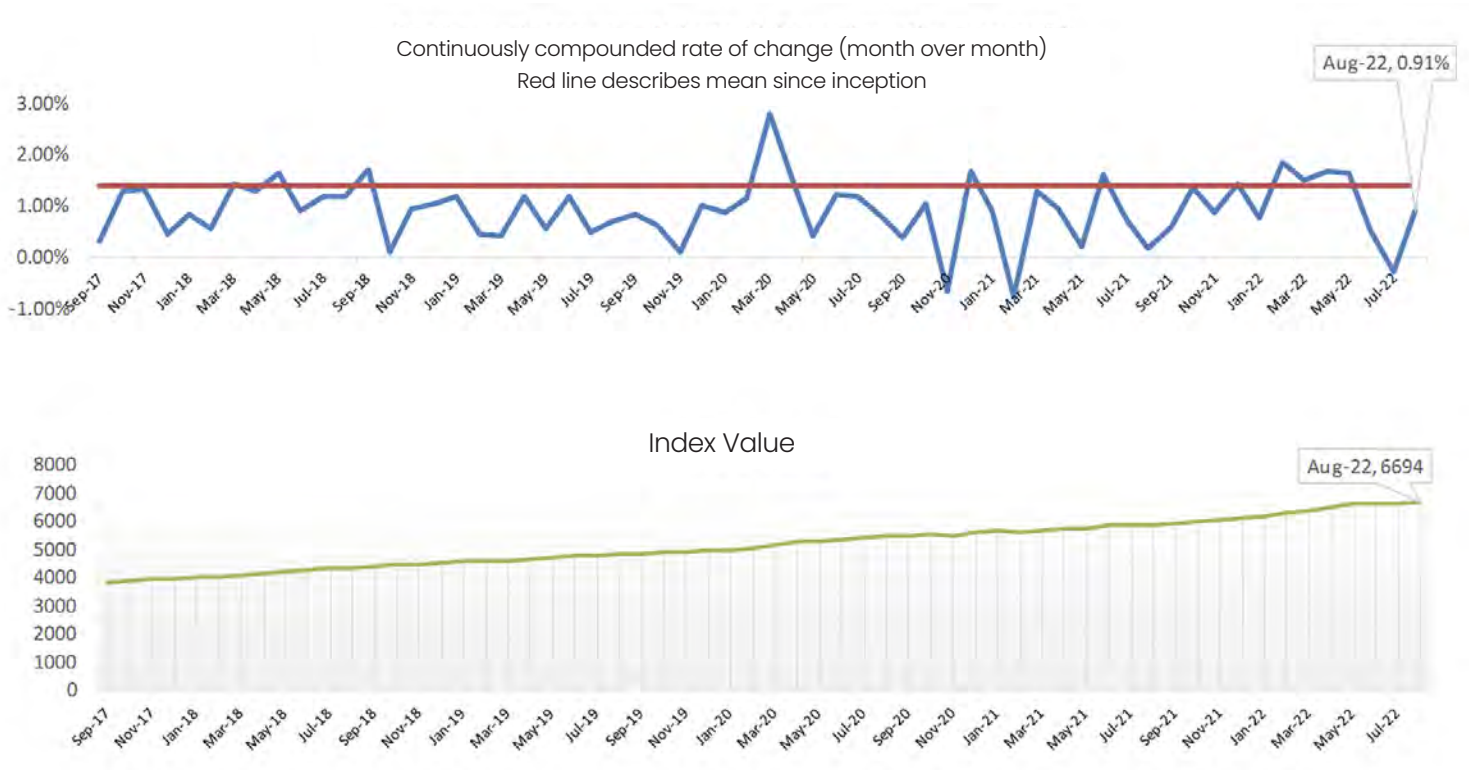

Figure 1: NYU ICS August 2022

EXPERIENCE TO DATE

In our work to investigate meaningful correlations to cyber events, our current findings include the following:

- **Continuous Posture Degradation.** Since its inception, the aggregate ICS value has been continuously increasing. This suggests that the experts agree that overall cybersecurity posture has continuously degraded over time. Given the constant evolution of threats and threat actors, this seems reasonable.
- **No Advance Incident Warnings.** No evidence could be found in the ICS sentiment data that advance indications were available to predict future consequential cybersecurity incidents such as specific data breaches. This suggests that the ICS is not useful as a predictive tool for decision making about day-to-day risk management, although it may be useful as a source of what the broader community is seeing.
- **Weight Manipulation Can Produce Correlation.** The percentages assigned to the five-level answers follow a balanced, symmetric scale between -20% and +20%, with smooth weighting intervals between the different answers. While these weightings did not produce meaningful correlations to publicly available security incident or vulnerability data, manipulations to answer weightings did allow for after-the-fact force-fitting of the ICS to support more visible correlations. By post-manipulating the weightings to force-fit a correlative match with previous, publicly available data, the case can be made that the weightings are calibrated to these past incidents and could thus produce more meaningful future guidance.

FUTURE DIRECTION

Our plans for the ICS include:

- Expanding the base of qualified responders.
- Investigating further to determine if weighting adjustments should be made to drive more meaningful correlations to cyber events.
- Building a more comprehensive attack/incident database to support our correlation efforts.
- Investigating the possible need for redesigning ICS questions.

For a deeper discussion on the ICS calculations, associated research and findings to date, please download **Experiences with a Sentiment-Based Cyber Security Index**. If you are interested in participating in the survey or have comments on this work, please contact me at andy.mccool@cybersecurityindex.org.

INTERVIEWS

AN INTERVIEW WITH EREZ ANTEBI, CEO, ALLOT

ALLOT: ENABLING SERVICE PROVIDERS TO SECURE CONSUMERS, FAMILIES AND SMALL BUSINESSES

Large enterprise teams have a massive advantage in cybersecurity—namely, their deep technical resources, trained security teams and focused attention from commercial vendors. This hasn't exactly solved the cyberthreat problem, but it has definitely helped CISO-led teams to improve their defenses every day.

Cybersecurity company Allot seeks to level the playing field for the rest of us. Their solution package for communication service providers can be integrated into existing infrastructure, extending security services to their consumer and small business customers. We recently chatted with the Allot team about their platform, and the prospects are promising.

TAG Cyber: What types of threat challenges do consumers and small business experience?

ALLOT: Consumers and small businesses both face threats from cybercriminals, though somewhat different in scope. Consumers expect to receive their internet connectivity similarly to the way they receive other utilities—ready to use. For example, when you get your water from the water company, you expect to be able to open the faucets and get clean water, without having to filter or condition it. Today, though, the responsibility for clean connectivity, free from cyberthreats, generally falls into the lap of the consumer. If a consumer wants safe connectivity, they generally need to download, install and maintain an endpoint solution—and only a small proportion of consumers actually do all of this properly, if at all. This leaves them vulnerable to viruses and other types of malware, adware and a growing number of phishing and banking scams that aim to steal the consumer's identity, data and money. Consumers are also increasingly open to attacks through their Internet of Things (IoT) devices, which are notoriously hard or impossible to secure directly. IoT devices act as open doorways to home (and business) networks, letting in botnet and other types of attacks.

Small businesses face all the same types of vulnerabilities as consumers. But they also have the added threat of work-from-home and bring your own device (BYOD) vulnerabilities. When employees use their own devices on the business

If a consumer wants safe connectivity, they generally need to download, install and maintain an endpoint solution—and only a small proportion of consumers actually do all of this.

network, or use business devices on their home networks, the protection solutions that might be in place in the business network can be bypassed by bad actors. Businesses are also more likely to be targeted by schemes to take their money. The most common, and growing, threat to businesses is ransomware, which holds the enterprise's critical data for ransom until they are paid to release it. Often, even paying high ransoms does not lead to the recovery of a company's data. There are network-based solutions, however, that can protect both consumers and small businesses without the hassle of installation and maintenance.

TAG Cyber: What are the major components of the Allot platform?

ALLOT: Allot Secure is a network-based cybersecurity platform offered to network operators. It protects consumers and small businesses through the operator's network, where all their data flows through already. Allot Secure gives service providers a way to include secure connectivity as a part of their core offering. At the heart of Allot Secure is NetworkSecure, which sits in the operator's network and protects mobile customers connected to the network. It blocks viruses, malware, phishing and ransomware attacks, as well as providing parental controls so that children can be protected from potentially detrimental content.

For the home network, there is HomeSecure, which also sits in the network and adds a thin client to the customer's home CPE. It automatically IDs all of the devices connected to the customer premises equipment (CPE) and protects them from cyberthreats, while offering parental controls, as well. When both NetworkSecure and HomeSecure are employed, there is a unified monitoring interface that also delivers consistent settings for each device and user when connected to the operator and home networks.

Finally, BusinessSecure is designed for small businesses. It is similar to HomeSecure in that it employs a thin client in the CPE. It offers group policy management, reporting tools, content filtering by category and other features that small businesses need for a simple, yet comprehensive, cybersecurity solution. Both BusinessSecure and HomeSecure rely on a cloud-based threat intelligence database that is frequently updated and refreshes the information on the CPE client to ensure up-to-date cybersecurity coverage. We also offer DNS Secure, which is a simple, transparent, network-based zero-touch cybersecurity solution for network operators who prefer a DNS-based solution.

TAG Cyber: How does your solution integrate with service provider infrastructure?

ALLOT: The solution is deployed on infrastructure owned by the service provider, which can be on premises, collocated or in the public cloud, and provides high customer data protection and privacy. This flexibility is enabled by the cloud-native technologies employed to build the Allot Secure platform from the ground up. The service provider interacts with this platform using a REST interface for service provisioning and, if functionality is integrated with the existing customer care portal, for configuration and reporting, too. From there, application of the subscriber-defined filtering configuration is done by the selected solution modules. If NetworkSecure is used, this will be done by the filters deployed on POPs. When HomeSecure/BusinessSecure is used in addition (or instead), it is done on the subscriber CPEs through a lightweight agent.

TAG Cyber: Tell us more about how service providers can grow revenue via a partnership with Allot.

ALLOT: We provide a combination of Allot Secure solutions, from which the provider can decide to implement. When the provider rolls out services based on these solutions, they are able to charge recurring monthly fees—usually between one to five dollars, depending on the market and service. With the higher-than-average take-up rate of these services, service providers can earn more revenue than with many of their value-added service (VAS) programs.

TAG Cyber: Can you share some insights into the future of cyberthreats in the upcoming years?

ALLOT: Attacks will be more frequent and more complex. This is the trend we have seen over the last few years through our research with existing customers. With 5G coming online, we are sure to see new types of attacks based in AR, VR and cloud gaming environments, as well as other new services enabled by 5G. At Allot, we are confident that we will continue to keep up with new cyberthreats, as we do today.

AN INTERVIEW WITH DEBBIE GORDON, FOUNDER & CEO, CLOUD RANGE

USING REALISTIC RANGE SIMULATIONS FOR CYBER READINESS WITH CLOUD RANGE

Security operations center (SOC) teams are often presented with unique situations featuring unknown threats that require fast thinking, efficient coordination and the optimal use of available tools. As one would expect, these objectives require practice to get right—and such learning is best done in a simulated environment, rather than during a live attack.

Cloud Range offers an effective platform and set of services that allows SOC teams to engage in realistic scenario training. The results range from helping an organization better gauge its overall security posture to helping participants better understand how to work together. We recently caught up with Cloud Range and learned more about their offering.

TAG Cyber: What are the toughest operational challenges of the modern SOC team?

CLOUD RANGE: One challenge is that it's difficult for SOC teams to show value to the rest of the organization. Cyber roles don't always align with business outcomes. SOC teams can be seen as an expense, and they may not get much attention unless something goes wrong. That's one reason why teams depend on Cloud Range's cyber-readiness platform, which prepares SOC teams to defend against the next cyberattack. It also measurably reduces cyber risk, which benefits the entire organization, including customers, shareholders and other stakeholders. Another challenge is that many SOC teams are taking on both traditional cybersecurity and operational technology (OT) cybersecurity. They need to know how to speak the language of industrial control systems (ICS) and IoT devices, as well as detect and mitigate potential IT- and OT-related cyberattacks. That is easier said than done. There are different objectives: IT prioritizes confidentiality of data, while ICS prioritizes the availability and integrity of processes and systems—and they work with different technologies. We are excited to be the first to have complete, full-service OT environments in our cyber range where teams can learn how to work together and defend against OT and IT/OT cyberattack scenarios.

There is also a high level of stress in SOC teams, which impacts operations because it can lead to burnout and the loss of good employees. Analysts should be seen as part of an organization's security stack; therefore, leaders must invest in ensuring their people are performing most effectively. That includes a combination of

Another challenge is that many SOC teams are taking on both traditional cybersecurity and operational technology (OT) cybersecurity.

one-to-one meetings, individualized career paths, professional development options for continued growth, and more.

TAG Cyber: What are the major components of the Cloud Range offering?

CLOUD RANGE: We are focused on providing SOC and incident response (IR) teams with the prescriptive, programmatic, hands-on training necessary to be ready for new threats and bridge any gaps, while proactively improving cyber readiness and measurably reducing cyber risk. The product that really put us on the map is our **live-fire team simulation missions** for SOC and IR teams. These immersive cyberattack simulations enable security teams to work together as a group in specific roles in an emulated environment to defend against the latest threats. Our FlexRange missions can be led by live instructors, so teams get real-time guidance, and security leaders get actionable feedback. Our customizable environments include the same industry-leading security tools that our customers use every day. We are constantly updating our library of IT, OT and IT/OT cyberattack scenarios, which map to MITRE ATT&CK frameworks. Simulations include options for blue teams, red team vs blue team, capture the flag, and more.

Another major component is FlexLabs, which offers over 1,500 on-demand training lab exercises, certification prep SkillPaks, challenge labs and skill assessments. Every lab maps to the industry-standard NICE Workforce Framework's set of knowledge, skills and abilities. Learning plans and courses can be prescribed and tracked in our Performance Portal. Additionally, our new aptitude and candidate assessments optimize hiring and help address the cyber workforce shortage.

TAG Cyber: How does your solution help with recruiting and retention?

CLOUD RANGE: The first step is to make sure that people start off in the right seat with the potential to succeed. Not everyone is going to be good at penetration testing, for example. Different cyber roles use different parts of the brain. RightTrak Cyber Aptitude Assessments use various modules—including pattern recognition, spatial visualization and more—to take the guesswork out of hiring cyber talent. Hiring managers can focus on candidates with the best cognitive fit for the job. FastTrak Candidate Assessments give hiring managers details that go beyond resumes or certifications. Candidates are immersed into a role-based simulation, allowing the employer to see how they perform in a real, practical way. FastTrak also eliminates any bias in hiring, opening the door for more diversity, equity and inclusion in cybersecurity. These assessments also help security leaders know how to best onboard new hires and build

a training program for them to grow and move up the career ladder, which leads to greater retention.

Plus, our FlexRange missions and FlexLabs help with more than training: they improve the overall culture. Cloud Range experts work with security leaders to build customized team and individual training programs, including live-fire simulation missions and hands-on skill development courses. As our experts identify gaps or see opportunities for improvement, we communicate with security leaders and outline the next steps. When people get to experience the training and development they need and want, they are more productive and more likely to stay longer. It's a win-win.

TAG Cyber: Tell us more about how security operations staff benefit from engaging with Cloud Range.

CLOUD RANGE: Our training does more than train. Our one-of-a-kind cyber-readiness solution improves communication, collaboration, problem-solving and more. It gives security team members and leaders confidence that they are better prepared to handle the next cyberattack. And it's fun! Teams love to be on the range, working together to figure out how to best detect and mitigate threats. It's like being in an escape room, but digitally. They don't know what they are searching for at first, so they have to follow the clues and work out what is happening and how to respond. It's exhilarating. We also handle all the administrative work—the coordination, execution, and other details—so the program is not adding to the noise, but reducing it. Additionally, everything that happens is tracked and analyzed. Our reports demonstrate measurable improvement in a security team's ability to detect and respond to cyberattacks, which is what boards care about.

TAG Cyber: Can you share some insights into the future of cyberthreats in the upcoming years?

CLOUD RANGE: Cyberthreats are increasingly encompassing or directly targeting OT systems. Gartner predicts that by 2025, cyberattackers will have weaponized OT environments to successfully harm or kill people. It's crucial for organizations to proactively prepare for cyberthreats and reduce their cyber risk. We are dealing with more than personal or sensitive data being lost—these new threats could lead to loss of life.

AN INTERVIEW WITH DON DUET,
CEO, CHAIRMAN OF THE BOARD
AND CO-FOUNDER, CONCOURSE LABS

SECURITY AS CODE INNOVATIONS FOR CLOUD-NATIVE ENVIRONMENTS FROM CONCOURSE LABS

It would be tough to find a business, enterprise or agency that is not moving at least a portion of its critical applications to the cloud—and many businesses are moving their entire application portfolio to a cloud-native hosting environment. This transition brings both opportunities and cyber risks.

Concourse Labs offers a Security as Code solution that supports a technique known as Cloud Native Application Protection Platform (CNAPP). We wanted to learn more about how the company's solution automated CI/CD security, while supporting the need for risk mitigation through runtime.

TAG Cyber: What do you mean by Security as Code (Policy as Code)?

CONCOURSE LABS: To us, Security as Code (SaC) means the formalization of security and control objectives into a set of automated rules and logic. Automating the testing and enforcement of controls and security objectives is the only way that organizations can repeatedly and scalably ensure that their software doesn't create unwanted risks and exposures. SaC can act over and govern many different elements of a software system, including the configuration of cloud services, privileges and identities, as well as Infrastructure as Code (IaC), itself—which is increasingly used to provision and procure cloud services. We don't think of SaC as a specific programming language or technology. Given the diversity of cloud service providers, cloud services and IaC technologies, SaC must support, by design, a diverse set of provider-specific "reference definitions" to ensure that its logic is consistent and safe with the technology and version being evaluated. Like all software, SaC needs to be managed and maintained. There are multiple change vectors independent of anyone's own view of a logical SaC implementation, including change to the provider specification, as well as additional resources and/or attributes that can be employed in the design and implementation of a strong control. The cloud does not stand still but is constantly innovating.

The cloud does not stand still but is constantly innovating.

TAG Cyber: What are the major components of the Concourse Labs platform?

CONCOURSE LABS: We provide a full-service, SaaS-based platform to help with all aspects of the implementation of an enterprise SaC program. There are five major components to our platform. First, an enterprise governance role and access-based framework enables foundational concepts, including the segregation of duties and layers of defense, along with information and organization barrier management. Next, there is policy lifecycle management functionality centered around a control-based methodology for SaC. This functionality includes control frameworks, policy collections, and a user-friendly graphical, no-code authoring experience. Third, our robust policy library spans pre-authored control frameworks, policy content for all major cloud providers, well-architected and foundational advice from providers, and research best practices and opinionated policy from our own Policy Research team. Next, we offer agentless continuous policy evaluation for cloud runtime state policy and API-based integration into CI/CD and SCM platforms for build-time policy evaluation. Finally, there is full audit and evidentiary record keeping for all aspects of SaC content management and policy evaluation results.

TAG Cyber: How does your solution provide automated security during the CI/CD pipeline?

CONCOURSE LABS: Our solution can easily be added to CI/CD pipelines via direct API integration or the use of our command-line interface (CLI) tool. As IaC software is moved through the pipeline control gates, the Concourse CLI/API facilitates scanning IaC file content with the appropriate policy set. Scans are performed on the Concourse SaaS platform and results returned back to the pipeline, where detailed policy violations and remediation guidance are recorded in the output and an overall pass/fail status determination is made. In this manner, Concourse enables policy evaluations to be performed on an event basis within the control plane of the pipeline, while policy content is managed external to the pipeline. This design supports both the segregation of duties, as well as the independent control of the policy governing the software. This allows the control-focused functions and team to revise policies as needed (zero-day exploits, for example) without the complexity and cost of having to work through multiple applications and DevOps teams to manage control content.

TAG Cyber: Tell us more about how the solution extends to runtime protection.

CONCOURSE LABS: Our solution was built from the ground up, with the point of view that strong governance requires the ability to assert policy at all layers of the application lifecycle—from SCM through CI/CD to ultimately the runtime itself and the active state of a cloud. The ability to identify and detect issues in the runtime is very dependent on time to detect (TTD), as well as time to remediate (TTR). Concourse maintains and manages a complete cloud state history by calling the APIs that describe the cloud state across services, networking, and identities and permissions on a continuous, near real-time basis. This approach allows Concourse to remain agentless; we just need permission to access the allowed APIs. This supports the creation of a history of cloud state over which policy can be applied and issues detected with a very low TTD. Many governance products offered by the major CSPs support a state refresh every four-to-six hours, whereas with Concourse, the continuous polling model supports refresh of state content in minutes. Concourse also enables risk-and-control teams to see all cloud states, regardless of whether they expect service usage in their organization. This is very helpful in the reconciliation of controls—cloud runtime is the source of truth for cloud usage—to actual behavior and usage, as it supports observability into rogue or unexpected service consumption. Having the complete cloud state history pre-positioned means control and policy development can be more agile and delivered faster.

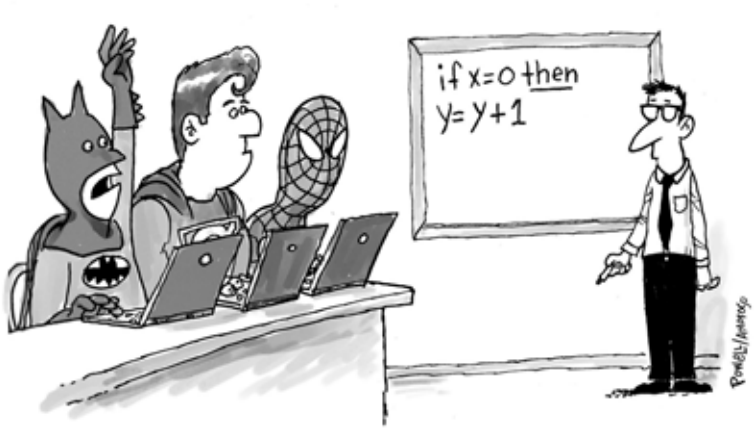

"So, workload containers are basically little Batcaves for apps?"

AN INTERVIEW WITH MASHA SEDOVA, PRESIDENT AND CO-FOUNDER, ELEVATE SECURITY

SOLUTIONS TO THE CHALLENGE OF INSIDER RISK MANAGEMENT FROM ELEVATE SECURITY

Ask any CISO what keeps them up at night, and it is likely you'll hear them mention insider risk. The challenge with insiders is that they are, by definition, trusted—and, as such, they can create problems by abusing authorized access. As a result, clever methods are required to address this subtle risk, especially in highly consequential threat environments.

Elevate Security provides a risk management solution for insider risk that is consistent with modern enterprise architecture, including cloud services, SaaS and third parties. We caught up with the Elevate Security team recently to better understand how their platform and technology can be used to reduce this risk.

TAG Cyber: What types of security issues emerge from trusted insiders?

ELEVATE SECURITY: Let's first define our terms. By insider risk, we are referring to the potential of any trusted worker to act in a manner that might harm the business—intentionally or, more likely, unintentionally. Insider threat refers to the actions of a specific individual, almost always one with bad intent. Not every risky insider becomes a threat, but every insider threat—malicious or accidental—starts with a risky insider. Attackers use a variety of techniques to victimize workers and create threats: phishing/smishing, infected websites, social engineering and more. We now know from Verizon's Data Breach Investigations Report (DBIR) and other research that 80% or more of incidents involve abuse of the human element. As recent headlines underscore, this relentless attack on our workforce is resulting in the loss or corruption of source code and other sensitive data, as well as ransomware, business outages, and the dilution of public trust and brand value. It's never been more important to identify at-risk workers and protect them from these attacks.

TAG Cyber: What are the major components of the Elevate Security platform?

ELEVATE SECURITY: Our platform identifies each worker's risk level and enables control technology throughout the security stack to dynamically adapt to each individual's level of risk. Unauthorized or high-risk activities are flagged, workers and management notified, and a spectrum of appropriate remediation enacted. High-risk workers may receive feedback and

80% or more of incidents involve abuse of the human element.

direction, or face additional authentication, authorization or technology controls that might be inappropriate or ineffective to apply in a one-size-fits-all manner. We also offer a rapid assessment service that identifies and scores all employees according to the current risk levels. This is the starting point for gauging insider risk levels and specific vulnerabilities. If clients wish to observe and report on risk over time, we offer risk monitoring on an annual basis. For clients ready to implement a predictive, proactive worker risk management program, we add feedback, rules, and the ability to deliver dynamic trust solutions based on constantly updated risk scoring. Another important offering, Elevate Intelligence, offers direct access to our risk database containing billions of data points, which clients access to determine attack trends, perform peer benchmarking, and increase the effectiveness and efficiency of cyber insurance—a key component of any risk mitigation strategy.

TAG Cyber: How does your solution integrate with other components of an enterprise's security infrastructure?

ELEVATE SECURITY: Our solution continuously gathers data from existing systems via API, and then normalizes and analyzes that data to define critical risks. It also constantly updates credit risk-like scores for each individual. This deep understanding of individual risk is then fed back to existing controls, enabling them to dynamically adapt to each person's unique situation.

TAG Cyber: Tell us more about how benefits emerge for both executive management teams and day-to-day practitioners when using the Elevate Security platform.

ELEVATE SECURITY: An Elevate-enabled dynamic trust model allows security teams to maintain worker productivity while protecting the business to the fullest. High-risk workers are identified and supported by reducing their risk levels and eliminating soft attack targets across the organization. We provide the executive team and board of directors with precise and easily understood metrics for internal workforce risk, both current and historical. This clear visibility allows executives to evaluate risk mitigation effectiveness, align the organization, drive accountability, and integrate workforce risk into ongoing governance, risk management and compliance (GRC) initiatives. Day-to-day practitioners can now apply risk-based controls to better protect their business, reduce the sheer number of insider incidents requiring investigation, and respond to incidents faster and with more information, all while better supporting the needs of their business partners. Insurers and brokers benefit from rapid risk assessments during the underwriting and renewal process, as well as gaining a deep understanding of risk trends for business planning.

TAG Cyber: Can you share some insights into the future of cyberthreats in the upcoming years?

ELEVATE SECURITY: The model of commodity attack tools is with us now, and likely for good. This weaponization has led to an increased sophistication and intensity of attacks on workers, which will continue. The novel application of commodity tools for achieving persistence is something we should all be watching for. When the barrier of sophistication for long-term adversary persistence is low, it will be crucial to manage trust and security controls in a dynamic way that adapts to the risk we can observe.

"My zero trust claims were more marketing than lies, really."

AN INTERVIEW WITH GRANT WERNICK, CO-FOUNDER AND CEO, FLETCH

GET AHEAD OF THREATS WITHOUT THE GRUNTWORK WITH FLETCH

Today's cybersecurity teams face a torrent of potential threats just about every second of the day. Global corporations with deep pockets and massive teams have a hard time keeping up with everything coming at them, while smaller companies face a much steeper climb.

Cybersecurity start-up Fletch was founded to solve the pains of these underserved teams. The company decodes cybersecurity analytics for clients ranging from small and medium-sized businesses to global corporations, in order to solve many of the most common problems encountered by security teams. We recently had the opportunity to interview the Fletch management team about their approach.

TAG Cyber: What are the types of questions and issues that tend to emerge for security teams as they do their work?

FLETCH: One of the biggest issues facing cyber teams today is the overwhelming number of threats coming at them on any given day. Thousands of new threats pop up, and thousands more evolve into something new. They need a fast, reliable way to track threats, figure out which ones matter to them and which ones don't. Knowing whether they are vulnerable to—or already showing indicators of—compromise from thousands of threats as they emerge is an impossible task for a company to do on its own. This is where Fletch comes in. We track and analyze over 30,000 security threat reports and articles each day and send alerts on the ones that matter to an organization—especially the ones that are likely to bubble up to the executive leadership level. We then provide a list of threats needing remediation, which is prioritized by severity and depth of impact. We gather all the info needed to quickly mitigate in one place, freeing up time to focus on what matters to the business.

TAG Cyber: What are the major components of the Fletch platform?

FLETCH: Everything Fletch does revolves around our natural language (NL) engine. It's an AI assistant that does dozens of tasks no human has the time or patience to accomplish. For example, Fletch's NL engine reads the backwaters of the internet every day. It extracts key information to figure out what threats matter, which articles have useful remediation advice, and the key indicators. It then generates new insights, including a severity score that determines the criticality and stage of a given threat, based on the sentiment of the cybersecurity community. It also goes the extra mile to automatically correlate its findings with a given

Knowing whether they are vulnerable to compromise from thousands of threats as they emerge is an impossible task for a company to do on its own.

customer's vulnerability and EDR products to generate personalized summaries, delivered right to their inbox, daily. Fletch already integrates with vulnerability scanners and endpoint products, such as Microsoft Defender, SentinelOne, Crowdstrike, Github Dependabot, Snyk, Qualys and VMware CarbonBlack—and more integrations are on the way. It would take a person days to do this process for one threat. Fletch is doing this for hundreds of threats discovered in thousands of articles every day.

TAG Cyber: How does your solution integrate with other aspects of the security architecture?

FLETCH: We're building a new layer for knowledge. We're helping companies free their best people from data plumbing. Fletch has built-in integrations to the most popular app and source code vulnerability scanners and EDR products. There's no custom data wrangling needed. Customers typically connect their products in less than 10 minutes via read-only APIs. From there, Fletch automatically correlates security indicators from a company's connected tools with the indicators of major threats, so the enterprise knows which threats are impacting them every day.

TAG Cyber: Tell us more about how you generate actionable insights for security teams.

FLETCH: Crawling the backwaters of the internet to find threats, determining which threats to focus on, and then correlating each threat's indicators with those reported by a company's tools requires a major investment of resources and people. Most teams can't, so they're left to cherry pick a very small set of threats to investigate, leaving themselves vulnerable to the massive amount of threats they can't investigate. Whether they are a security operator in a dedicated team, an army of one, or someone only focusing on security part-time, people don't have the time to run and ground every single burning threat emerging and evolving each and every day. They especially don't have time to play fetch, digging into every threat an exec reads about in the news. Fletch makes all this busy work disappear, while helping teams get days to weeks ahead of major threats. Many are calling it their daily information hub, and others joke that it's a threat intel team in their back pocket.

FLETCH: ***TAG Cyber: Can you share some insights into the future of cyberthreats in the upcoming years?***

The attack surface is ever expanding. It used to be adversaries focused on systems, and now it's the software supply chain. It's just starting to be the SaaS applications we all use every day. No one human or team can keep up. Most can't even hire the right people to do the right processes to defend themselves. The lens on the problem needs to change. AI can handle the data-munging gruntwork. Our easy-to-understand dashboards and notifications simplify the process, and we then layer on expert advice to help speed remediation regardless of a team's capabilities.

AN INTERVIEW WITH ANDREW HOYEN, PRESIDENT & COO, INFINITE GROUP, INC. (IGI)

ADVANCED CYBERSOLUTIONS FOR SMALL TO MEDIUM-SIZED ENTERPRISES FROM IGI

One aspect of the modern enterprise that all experts agree upon is the high level of complexity that must be addressed by security teams. New techniques are required to address this complexity so businesses can monitor networks, address weaknesses and take corrective action.

IGI works closely with customers and partners, including MSPs and MSSPs, to help reduce risk through their service portfolio and SaaS-based product offerings, including their patented Nodeware solution. IGI has revolutionized the vulnerability management market for SMBs through flexible, lightweight deployment combined with continuous scanning to significantly reduce risk and increase visibility for companies and partners alike. We were excited to get an update on the work at IGI and how they continue to provide effective cybersecurity solutions.

TAG Cyber: What are the primary cyber risks that emerge from the complexities of the modern enterprise?

IGI: The work-from-anywhere trend exposes security gaps for organizations, and these gaps need to be quickly addressed. Attackers are constantly evolving and developing new ways to penetrate networks and systems, and companies need to stay at the forefront of security technology to thwart these attacks. Our Nodeware solution serves that purpose, as a first line of defense for many companies to understand where these exposure points lie and how easy they are to exploit. The emergence of the cloud has also lulled people into a false sense of security. In fact, if your cloud isn't set up properly and securely, it increases the number of access points into the enterprise. The greater the attack surface, the greater the exposure to the company.

Additionally, companies are constantly bombarded with the latest tech solutions, but don't have the ability to pull them all together to work effectively. With resources stretched, and companies not having a plan, many decisions are made without proper guidance. Regulations at the federal and state levels have caused a great deal of confusion and uncertainty for the enterprise. Due to increasing activity by threat actors, companies need to be hyper-diligent and vigilant in their approach to protecting their most valuable asset: data. Moreover, when it comes to today's job market, it is increasingly difficult to attract and retain talent in cybersecurity.

The emergence of the cloud has also lulled people into a false sense of security.

Combine that with the cost of this expertise, and enterprises are left wondering how and where to address their budgets. As a result, companies are outnumbered and look to third-party organizations, like IGI and their partners, to support them effectively. IGI and its partners bring support to the most critical areas of the enterprise through a combination of human involvement and technology.

TAG Cyber: Tell us about IGI and the work you do.

IGI: We are a cybersecurity technology and services innovator, committed to helping clients and partners solve the complex challenges related to safeguarding themselves against outside threats. We are the original equipment manufacturer (OEM) of the rapidly growing, patented Nodeware vulnerability management solution that provides continuous vulnerability scanning, thereby helping organizations lower their risk of a cyber attack. This game-changing, cloud-based SaaS solution has enabled our growing partner base to integrate Nodeware into the stack of solutions they bring to their end customers. We also have a full suite of cybersecurity services, including a new twist on the traditional virtual CISO (vCISO), aptly named CISOTaaS, where we bring a team—not just one person—to support your needs, along with risk assessments, penetration testing and incident response. Our approach is simple. We spend the time learning about our clients' and partners' businesses to provide them with personalized solutions that align with their business strategy, while helping them strengthen and create more resilient cyber defenses.

TAG Cyber: How does Nodeware integrate with existing security infrastructure?

IGI: Our advanced API enables partners to integrate Nodeware into their stack of solutions and feed the data into their single pane of glass. We also offer a fully functional portal application to display key vulnerability data that both the partner and the customer can interact with. Nodeware can be deployed via a network software sensor and as an agent to capture both network and cloud-based assets. It provides our customers and partners with continuous visibility into both internal and external facing assets, while uniquely being able to see anything that occupies an IP address, whether it be Windows, Linux, Mac or IoT, to name a few. We have built a flexible, lightweight, zero network degradation solution that identifies all assets and then runs critical vulnerability scans on a continuous basis to provide real-time information.

TAG Cyber: Tell us more about how you help customers optimize their vulnerability management process.

IGI: The first step in vulnerability management (and cybersecurity) is knowing what you have, or what is connected

to your network. Nodeware monitors the network on a continuous basis, creating a dynamic inventory of the customer's assets and providing the most up-to-date vulnerability information for those assets all in one place. In addition, Nodeware attaches a vulnerability risk ranking to each asset, which classifies vulnerabilities based on their criticality, and helps customers prioritize their remediation efforts. Given new compliance regulations around continuous scanning, Nodeware provides a solution that has brought the market to us, whereas other solutions are still point-in-time scanning solutions.

TAG Cyber: Can you share some insights into the future of enterprise security in the upcoming years?

IGI: We will continue to see the consolidation of cybersecurity companies in this market, not only at the enterprise level, but also downmarket. Cyber and technology companies will continue to chase the "holy grail" of one solution fits all, in order to make it easy to manage the attack surface, but that is not reality. It takes a combination of the right technology with the right people to manage this effectively. Ultimately, no one is completely safe and secure, but you can take the necessary steps to follow good cyberhygiene to stay ahead. Unfortunately, regardless of size, everyone is a target. While enterprise is still a big target, the SMB/SME market is now a focus for threat actors to drive disruption and profiteering. It's no longer the hooded sweatshirt hacker in his basement; it has become nation-state backed exploitation from places like China, North Korea, Russia and more. Lastly, cyberliability will be an area to pay attention to, as more companies are impacted by these events. We are seeing more claims and the subsequent increase in cost for insurance. It's critical that enterprises have plans in place to address their cyber needs as the underwriters are paying closer attention. With the C-suite and company boards of directors now affected, liability has an even bigger impact on the enterprise as a whole.

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.

"Next missing tooth, I'm asking for Bitcoin."

AN INTERVIEW WITH AJ SARKAR,
FOUNDER & CEO, OPTIMEYES.AI

THE OPTIMEYES AI-POWERED RISK MANAGEMENT FRAMEWORK

Building a risk-based strategy in modern business is a particular challenge, especially in the context of evolving cyber and data privacy threats. The goal of such a strategy is to use impact assessments and visual analyses to drive decision-making. This is best done in the context of understanding an organization's risk tolerance, perception of threat, compliance requirements, and integration of non-cyber-related risks.

OptimEyes offers a SaaS platform that supports the quantification of cyber and privacy risks through the use of advanced processing methods based on artificial intelligence. We spent time recently with the OptimEyes team to learn more about how such methods can be used to help clients on an effective risk-quantification journey.

TAG Cyber: What types of cyber and privacy risks does your solution address?

OPTIMEYES: Our Integrated Risk Modeling and Decision platform helps business leaders address a wide range of cyber and data privacy risks. Risks include ransomware, IP theft, data loss, insider threats, regulatory and compliance failures, product flaws, supply chain compromise, cross-border data transfer, data privacy disclosures, and health and safety concerns. As there is often a causal link between a cyberattack and a data privacy breach, we quantify, aggregate and integrate these risks in our platform to create a single source of truth. These analytics support the different needs of multiple stakeholders, including Chief Information Officers, Chief Information Security Officers, Chief Privacy Officers and Chief Risk Officers.

TAG Cyber: What are the major components of the OptimEyes platform?

OPTIMEYES: We listened to approximately 1,000 CXOs when building our platform. We use a Zero Trust risk management approach with integrated risk modeling and decision-making capabilities. It is critical to link the impact of risk on an organization's business goals and objectives, and this understanding is central to our approach. Companies are looking for smart ways to pivot their traditional, governance-based risk management programs to a vision that is forward looking, works in real time and leverages opportunities provided by today's technology. Realistically, to get ahead of today's cyber and data privacy challenges, the old way of managing risk just isn't good enough. You don't see hackers and other bad actors using

To get ahead of today's cyber and data privacy challenges, the old way of managing risk just isn't good enough.

slides and spreadsheets! Our platform offers a unique approach and has numerous components, including inside-out modeling (i.e., the utilization of an organization's own data rather than data from industry averages), risk quantification and industry benchmarking to generate actual dollar numbers rather than typical red, amber and green risk descriptions. This creates meaningful risk exposure and risk remediation cost analytics. Our powerful AI/ML drives risk scenario planning by creating what-if scenarios for leadership decision-making.

Recognizing that organizations are required to comply with multiple regulations and follow the standards of many risk frameworks, our Unified Common Control Framework identifies and maps the commonality of those regulations to create common controls for enterprise-wide risk quantification, as well as the understanding of control effectiveness. In addition to providing assurance that local requirements are being followed, the user benefits from a broader, more holistic understanding of risk. The platform is easy to use, provides multiple on-demand use cases and can be configured very quickly with organization-specific information. Coupled with our persona-based intuitive, neuroscience dashboards, our aforementioned capabilities provide a powerful platform for key stakeholders to prioritize risk-mitigation steps and decide how to deploy investment dollars—which is the heart of effective decision-making.

TAG Cyber: How does your solution use artificial intelligence?

OPTIMEYES: Our platform uses AI/ML to help organizations reduce risk exposure, annual loss expectancy and optimize risk modeling functionality, as well as make the platform easy to use, operate and maintain. Platform functionality includes built-in risk scoring with risk exposure and remediation cost quantification and optimization. This core AI/ML functionality supports industry risk benchmarking, risk scenario planning, the operation of our aforementioned Unified Common Control Framework, and the detection of risk score anomalies. Combined with machine learning and the further utilization of any client and industry data already ingested and stored by the platform, OptimEyes supports an organization's efforts to predict risk and improve forward-looking, risk-mitigation and investment decisions.

TAG Cyber: Tell us more about how clients can quantify cyber and privacy risks with additional non-cyber-related risks to their business.

OPTIMEYES: As described above, our platform helps an organization understand the impact of cyber and data privacy risks on their business goals and objectives. Our models map cyber and data privacy risk to multiple risks that can impact a

business, including business continuity, competition, geopolitical, people/talent, product development, supply chain, regulatory compliance and enterprise operations. An organization's business risk profile is, of course, continually evolving. The platform can be updated in real time to reflect emerging business risks deemed important by an organization.

TAG Cyber: Can you share some insights into the future of cyberthreats in the upcoming years?

OPTIMEYES: Cyberthreats are about to change dramatically. Whereas today, companies struggle to understand the risks and remedies from threats that can harm profitability and privacy (e.g., theft of intellectual property, compromise of customer addresses), tomorrow's threats will harm our personal well-being and property. Think about it: The technology marvels of the next ten years are the ones quickly inhabiting our physical environment—our kitchens, conference rooms, cars, production lines, personal medical devices and pollution remedies. The upside of these technologies is profound: improved quality of life, economic opportunity and increased health and safety. But they also create new threats. Not only can a malicious person from "traditional" cyberspace take our photos and financial records encoded with the cyberdomain's 1s and 0s, but they will now also be able to manipulate or steal the 1s and 0s of our physical domain to do things such as: peep inside rooms and offices considered private; cause "accidents" by delaying train or car brakes; create severe economic hardship to both big and small companies by causing "inexplicable" failures in manufacturing lines; change medical information on wearable healthcare devices; and leak highly personal information about a person's well-being. Understanding the potential risks—both in traditional cyber and emerging physical domains—has become more urgent than before and more personal than ever.

Responsible company managers must use risk-modeling tools to ensure that their business goals are underpinned by robust and safe technologies. Moreover, those designing and operating these technologies will need such tools to ensure that their security and safety strategies are effective despite their underlying complexity. We live in exciting times. There is much opportunity ahead for greater efficiency and effectiveness in business, as well as improved quality of life at home. But with great opportunity comes great responsibility, which makes the modeling and management of technology-based risks more urgent than ever.

AN INTERVIEW WITH BOAZ AVIGAG, SENIOR DIRECTOR OF PRODUCT MARKETING, PERIMETER 81

SECURING YOUR CORPORATE NETWORK WITH PERIMETER 81

Common challenges in a corporate networking environment include lengthy deployment cycles, complex management processes and poor visibility into network-connected resources. This requires an improved means to build, manage and monitor a modern enterprise data networking infrastructure.

Perimeter 81 offers a simple platform to set up, configure and manage a network. The monitoring and security functions are integrated, so that network administration becomes possible from a single dashboard. We asked the Perimeter 81 team to share insights into how their customers can benefit from such a simple solution.

TAG Cyber: What are the current challenges enterprise teams experience as they try to manage their network?

PERIMETER 81: Enterprise teams are looking for ways to help them adapt to the new enterprise network topology. Instead of static branch locations connected to static data centers, the new enterprise network consists of users who connect from a wide range of locations: namely the office, home, coffee shop, airport and pretty much anywhere else. Additionally, applications have been, at least in part, pushed out to cloud environments, such as public clouds and SaaS. The traditional enterprise network has essentially dissolved and been replaced by the internet itself. This raises a wide range of both networking and security challenges, since the need to securely connect any possible path between a work-from-anywhere user to a delivered-from-anywhere application raises several requirements. Firstly, it requires a more scalable and secure solution than traditional VPNs can offer. These were not designed to cope with so much of the modern hybrid workforce connecting remotely, so much of the time. Secondly, it requires an optimized routing scheme that doesn't backhaul all of the enterprise's traffic through a static location, such as an on-premises data center. Thirdly, it requires security capabilities that can manage secure access to different cloud environments, as well as protect employee web browsing to avoid the downloading and distribution of malware within the network.

The traditional enterprise network has essentially dissolved and been replaced by the internet itself.

TAG Cyber: What are the major components of the Perimeter 81 platform?

PERIMETER 81: First, there is the platform itself, which consists of more that 40 globally distributed Points of Presence (PoPs). The platform provides services such as SSL inspection, identity provider (IdP) integration, monitoring and analytics, SIEM integration, and others. Two main services run on top of the platform: Zero Trust Network Access (ZTNA) and Secure Web Gateway (SWG).

ZTNA enables employees to securely connect to the enterprise network and provides granular access control to all enterprise resources. Employees run a dedicated agent on their endpoint devices to connect to the ZTNA service. Agentless access is also available to support secure connectivity for unmanaged devices. SWG secures employee web browsing by limiting access to harmful or restricted websites, as well as by scanning downloaded files for potential malicious malware and blocking it when detected.

TAG Cyber: How does your solution support rapid provisioning and deployment?

PERIMETER 81: The greatest advantage of a cloud-based network—also known as Network as a Services (NaaS)—is that there is no need to deploy anything. There is no hardware involved; the solution is fully delivered from the cloud, meaning it is already deployed and running, waiting to onboard customers immediately. The provisioning process is greatly shortened since all configurations are performed in a single-pane-of-glass management console. Configuration changes are propagated to all PoPs automatically, eliminating the need to configure point solutions in multiple locations. The agent download and set-up process is fast and requires minimal user intervention. A typical on-boarding process can take as little as 15 minutes to complete, and rarely takes more than a few hours—even for the most complex network topologies.

TAG Cyber: Tell us more about how your platform enables secure Zero Trust access for customers.

PERIMETER 81: Zero Trust is based on a few principles, the most being the concept of least privileges. This means that users are not granted access to any company resources by default. Even after the authentication stage, where a user's credentials are verified, preferably by a multi-factor authentication (MFA) scheme, no access to resources is granted. Only after a specific user, or a defined user group (e.g. R&D) are explicitly granted access to a specifically defined resource, will access be granted. Granular user and application-specific access rules are key to limiting lateral access and ensuring minimal risk in the case

that user credentials have been compromised. Another aspect of Zero Trust is to continuously monitor endpoint compliance using a Device Posture Check (DPC) mechanism to make sure all required security measures for the device are met throughout the session's lifecycle.

TAG Cyber: Can you share some insights into the future of cyberthreats in the upcoming years?

PERIMETER 81: Cybersecurity is essentially a cat and mouse game. As security solutions are evolving and improving, so is the sophistication of bad actors trying to overcome them. With the continuous digital transformation trend, an increasing share of the business world is becoming more digital, as well as more distributed in the cloud. When we combine this trend with the wide adoption of remote work models, we see how the new enterprise network constitutes a vast and virtually endless attack surface that bad actors will try to exploit. We will likely see a growth in social engineering techniques and attacks, ransomware and other monetizable attacks, and third-party supply chain attacks, which enable perpetrators to compromise potentially hundreds or even thousands of companies in a single breach. Enterprises that step up their game, adopt the right solutions and implement security best practices are much less likely to be breached. Those who don't will become sitting ducks for hackers trying to make a quick buck or realize alternative objectives. A typical enterprise can expect to be the target of hundreds, if not thousands, of attacks every month. Hoping for the best will likely not be a winning strategy. Beefing up network security likely will.

"I hardly recognize the firm since we got into cyber."

INTERVIEW WITH KARIM HIJAZI, CEO, PREVAILION

UNDERSTANDING COMPROMISE INTELLIGENCE FROM PREVAILION

As enterprise teams attempt to prevent attacks, the need arises for Compromise Intelligence, which incorporates insights from sources such as hacker networks. This intelligence helps defenders avoid issues that might be arising so that preventive or response actions can be taken to minimize negative consequences.

Prevailion provides world-class Compromise Intelligence capabilities for enterprise security teams. We wanted to learn from this industry leader about how their commercial solution can be used to reduce the risk of compromise and assist with critical tasks such as securing the enterprise supply chain.

TAG Cyber: What is Compromise Intelligence?

PREVAILION: Compromise Intelligence is egress communications gathered from adversary-controlled malware infections and replicated malware egress for security validation originating from within an enterprise.

TAG Cyber: How do Prevailion's products work?

PREVAILION: We have a vast sensor net of formerly malicious command-and-control endpoints that we've commandeered and repurposed to gather Compromise Intelligence on over 70,000 organizations globally. We can leverage this for compromise telemetry and to validate security controls.

TAG Cyber: How do you derive and obtain information from your intelligence sources?

PREVAILION: We identify and convict malicious command-and-control endpoints and passively collect proprietary compromise telemetry. The malware itself provides telemetry for enterprises with active infection. In the case of our Arktos product, our malware replication agent continuously tests security controls for failure due to infrastructure drift, configuration changes and threat TTP evolution.

TAG Cyber: Tell us more about how enterprises can deploy and integrate Prevailion solutions.

PREVAILION: Enterprises can use Compromise Intelligence to assess the historical success or failure of the organization's security. Additionally, the employment of continuous security-validation

technology will preemptively tune security controls against sophisticated threats.

TAG Cyber: Can you share some insights into the future of enterprise security including supply chain risk?

PREVAILION: We believe that forward-looking enterprises will prioritize Compromise Intelligence to assist in decision support, risk posture and security validation. This will allow organizations to focus on the urgent needs of active compromise, versus potential threats and vulnerabilities.

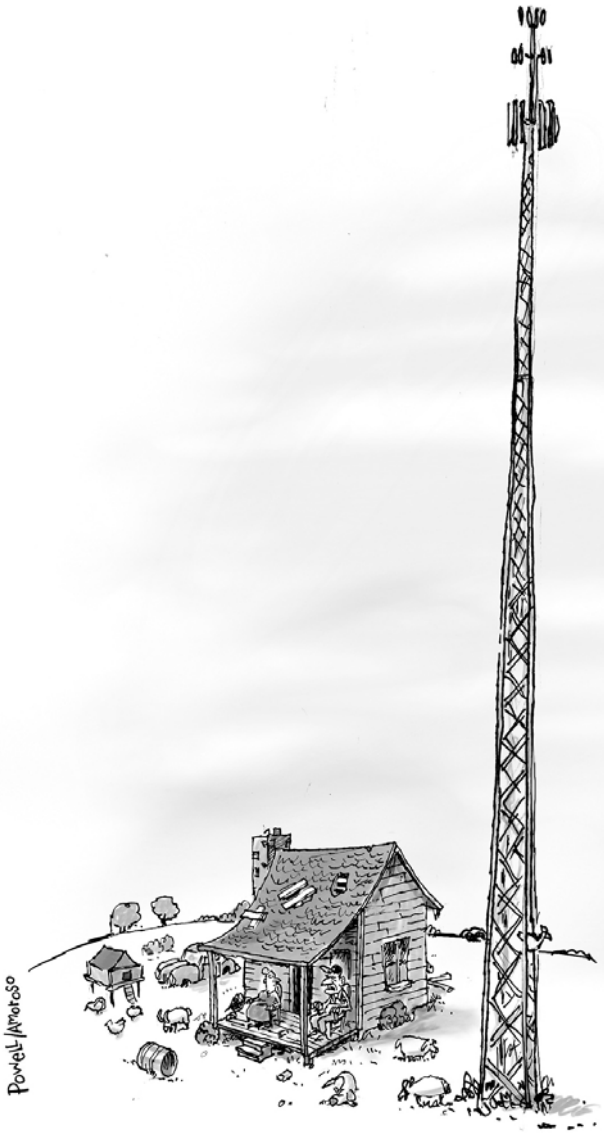

"Our Internet's been purty good since Goober started workin' for AT&T."

AN INTERVIEW WITH J.J. GUY,
CO-FOUNDER AND CEO, SEVCO SECURITY

CYBER-ASSET ATTACK SURFACE MANAGEMENT USING SEVCO SECURITY

Modern enterprise architectures with data and workloads hosted locally and in public cloud infrastructure have led to an increased need for more accurate asset inventory and security management. Dealing with asset risk across the modern enterprise requires a new approach to visibility, analysis and the identification of gaps that require security action.

Sevco Security provides a commercial platform for this task, which is referred to as cyber-asset attack surface management. We are excited to learn more from Sevco about how their solution can complement existing tools, improve incident response, and maintain security compliance.

TAG Cyber: What is the challenge to managing cyber-asset inventory?

SEVCO SECURITY: The challenge of managing cyber-asset inventory comes down to one main point: enterprises cannot protect what they cannot see. This foundational cybersecurity problem stems from the uncertainty of enterprise inventory—the attack surface—and the complexities around securing a growing number of IT assets. Unfortunately for today's modern enterprise, the way that IT tracks and secures assets is fundamentally flawed. Typically, we see every security program framework claim "an accurate inventory of enterprise assets" as a foundational control. However, it is increasingly apparent that these teams are failing to understand the true inventory of assets, no matter how often IT teams claim to have asset-inventory under control. This is why we founded Sevco in 2020—to fix a decades-old problem of visibility into asset inventory.

TAG Cyber: How does the Sevco platform work?

SEVCO SECURITY: We built Sevco to be a security asset intelligence platform that organizations can use as the foundation for their security programs. Sevco is a cloud-native platform that delivers converged asset inventory and generates real-time asset telemetry, and then publishes both for use by other IT systems. Our platform is a cloud-native solution that delivers a consolidated view of assets reconciled across all IT systems. More importantly, with asset telemetry, the platform has the ability to understand how assets change over time. Our unparalleled asset telemetry gives us the unique ability to deliver real-time asset inventory while simultaneously tracking the

With asset telemetry, the platform has the ability to understand how assets change over time.

state of every asset across the enterprise at any point in time. Our platform integrates with existing IT and security systems via native APIs to pull their specific view of asset information. Customers use our platform to make sense of the data they already have, which makes their existing security/IT products and procedures more effective with little to no change to them.

TAG Cyber: How does your solution support continuous compliance?

SEVCO SECURITY: Our solution makes reporting and maintaining compliance an easy and consistent task for enterprise IT and security teams. Compliance initiatives need to be built with a comprehensive understanding of the operational environment. Assets within the enterprise are highly dynamic, and compliance efforts are set to fail from the start without an accurate, real-time account of device inventory, types and operating conditions. Sevco makes monitoring, reporting and maintaining compliance easier and more consistent. The platform can aggregate multiple sources and duplicate inventory reports to provide a complete picture of IT assets. It also provides the ability to organize in-scope devices and separate the reporting against out-of-scope assets to ensure that compliance needs are met for every device. Lastly, Sevco's continuous processing provides IT and security teams with the ability to monitor build standards for in-scope systems.

TAG Cyber: Tell us more about how your solution supports attack surface management.

SEVCO SECURITY: Our Asset Intelligence Platform supports attack surface management by providing organizations with a single source of truth around the current IT assets that make up the attack surface. Ours is the only solution on the market to deliver real-time inventory and track the state of every asset across the enterprise at any point in time. The platform is driving innovation within the attack surface management market in three specific ways. First, Sevco delivers a converged, multi-source asset inventory. This inventory provides customers with a complete view of their assets. Another key innovation is our unique enterprise asset telemetry. This telemetry technology provides customers with knowledge around why inventory might change from one day to the next. And lastly, Sevco publishes data into existing IT and security systems, such as the CMDB, SIEM, and SOAR platforms, so that customers can act on better data with little change to existing processes. To reiterate, an organization cannot protect what they cannot see. We understand that the modern-day enterprise is constantly changing, and the assets you protect today may not be the same assets you protect tomorrow. For that reason, our platform is easily extensible to incorporate new sources (users, applications, cloud infrastructure,

etc.) that support additional asset classes, as well as support new applications that are developed by us, third parties, or customers.

TAG Cyber: Can you share some insights into the future of enterprise security in the upcoming years?

SEVCO SECURITY: I envision that many organizations will become overwhelmed with all the information, tools and new compliance processes that are supposed to help improve an IT security posture. Instead, I believe enterprise IT and security teams will take pause and revert back to the basics of security, by focusing on gaining a better understanding of the assets they currently have to inform the security protocols needed to protect sensitive data, devices and employees. Over the past few years, there have been lots of changes to our business environment—from the migration of IT environments to the cloud, to transitioning whole workforces to a work-from-anywhere structure—and there are more complexities around asset inventory and security than we've ever seen before. And the sad reality is, we never really had a good understanding of these IT assets to begin with. Before we tackle the next new security fad, I think we will see organizations re-establish security foundations. This will have the most significant impact on security programs and set them up for success in the future.

"Your daughter is just not grasping continuous GRC workflow."

AN INTERVIEW WITH ROSARIO MASTROGIACOMO, VP ENGINEERING, SPHERE

ESTABLISHING IDENTITY HYGIENE TO SECURE ENTERPRISE ASSETS WITH SPHERE

A major goal for any enterprise is to ensure that only desired individuals and groups are authorized to access information in applications and systems. This requires that proper identity hygiene be ensured throughout the access lifecycle, since securing assets can only be done with accurate and complete identity information.

SPHERE offers a platform that enables identity hygiene through the use of actionable intelligence from the infrastructure. We spent some time recently with the SPHERE team to learn more about identity hygiene, as well as how data sources are identified and remediation is achieved at scale.

TAG Cyber: What is identity hygiene and why is it important for enterprise?

SPHERE: Identity hygiene is a combination of activities that organizations perform to maintain the security of their data, infrastructure, and applications. It's a practice that people and enterprises just know they need to do—very similar to personal hygiene. We know we need to brush our teeth and shower on a regular basis because if we don't, our hygiene gets worse and ultimately, our whole body is at risk. Identity hygiene follows the same concept and that's what SPHERE delivers. It's a practice that companies have been trying to do for many, many years, but with different levels of success. If you don't address your identity hygiene, security issues and breaches are almost certain to happen. What SPHERE does is look at the human involved. We look past the account and focus on the identity to determine the risk. Take Bob Smith for example. It's not about Bob Smith's account, it's about the human, Bob Smith, and what he *actually* has access to versus what he should have access to. Databases, applications, SharePoint sites—we look at all these things to get a fundamental understanding of who Bob Smith is, who that identity is, and all the accounts Bob controls. This is truly the cornerstone of enterprise security in today's age. It's focusing on the human, whereas in the past, it was always just about the account. Fundamentally, you must go past the account and look at who is controlling it and the role that person has in the organization to understand and establish identity hygiene.

Control over access is moving away from the technology department and instead going directly to the end user.

TAG Cyber: What are the major components of the SPHERE platform?

SPHERE: The platform is made up of four major components: connectors, receivers, the business intelligence layer, and the user interface. The connectors go out and collect data from target systems, while receivers listen for activity data. Our BI engine is where all the magic happens. The engine takes the collected data, correlates it, and finds ownership and security issues. What we do with this raw data is the value-add we provide, the reason for the product, and why customers really, really love our products and services. Once we have processed all this raw data, we introduce the final component, the user interface. Through our ARM workflow tool, users can interact with data owners to move on to the next step: remediation. This is where users can remediate, solve issues, and achieve a state of clean, clear identity hygiene.

TAG Cyber: How does your solution integrate with identity and access management (IAM) and other related enterprise protection systems?

SPHERE: We integrate in two different ways. We pull data from a lot of IAM solutions, because it enhances our algorithms. We do this by leveraging a connector that pulls data to do additional analysis, which enhances our workflows. But we can also send data to systems, where we essentially clean the data, apply hygiene, then feed or send it to the IAM database. We call these IAM feeds or SPHEREfeeds.

TAG Cyber: Tell us more about how data is collected, analyzed and used to remediate issues.

SPHERE: Our agentless connectors and receivers are how we collect raw data. We don't do reporting for the sake of reporting. Our reporting is purposeful and provides paths for remediation. To that end, our connectors are also purposeful, they collect the bare minimum of raw data we need in order to accomplish this goal.

TAG Cyber: Can you share some insights into the future of cyberthreats in the upcoming years?

SPHERE: I think the biggest issues that people must grapple with involve two different trends: control over technology and how control over access is moving away from the technology department and instead going directly to the end user. Think SharePoint sites, Dropbox, OneDrive…all these technology platforms are going straight to the end user and allowing them to share data. In the past, people would have to go to their IT person to ask permission to share and collaborate with a colleague on a secure platform or be forced to email their colleague with an attachment. Now, we're moving in the direction where IT isn't involved, and an end user presses a

share button to collaborate without the IT team even knowing. This is great for an end user—no more roadblocks, no more extra steps in collaboration—but a nightmare for IT. Oversharing is a huge issue companies have to grapple with. We must get sophisticated enough to understand the difference between good sharing and bad sharing. Sharing and collaborating are necessary and, in some instances, critical to the success of a business, but "bad sharing" is a real security threat.

So why do we overshare? The more data, the better, right? There is the idea that advanced analytics and machine learning can add value to a worker. That extra data can be fuel for new, sophisticated algorithms and offerings for the organization that ultimately provide data for new sales or marketing opportunities that help the business grow. Sounds great. But that counters the idea of limiting data for security threats and giving employees the bare minimum of what they need to do their job. It's a battle, which means we need to find a balance. An appropriate amount of sharing, yet not enough to hinder or limit the data an individual can use and access, but also not threaten or risk the business. At the end of the day, employees are processing information to better the business based on the data they're able to access. This is how we mitigate future threats. The idea is to provide the data needed to find new opportunities, while avoiding insider threats and security issues that compromise security.

"I'm worried that our software will break when we go from BC to AD."

AN INTERVIEW WITH CODY CORNELL, CO-FOUNDER & CHIEF STRATEGY OFFICER, SWIMLANE

LOW-CODE SECURITY AUTOMATION FROM SWIMLANE

The security community has long desired an effective implementation that supports XDR, primarily to take advantage of automation in the SOC. Low-code security automation delivers the desired outcomes of XDR through a unique approach that helps security teams be more active, autonomous and adaptable to threats.

Cybersecurity vendor Swimlane provides Turbine, a commercially available automation platform that secures modern enterprise systems. The platform uses low-code techniques to remove the need for developers to build complex integration with other systems. It also extends visibility and actionability, as well as ultimately unifies workflows, telemetry and teams through a single system of record. We spent time recently with Swimlane to better understand how this works.

TAG Cyber: What is meant by low-code security automation?

SWIMLANE: Low-code security automation is the sweet spot between legacy SOAR and no-code automation solutions. It delivers the robust application development capabilities that you'd expect from SOAR, coupled with the ease of use and fast time-to-value advertised by no-code security automation tools. The beauty of low-code security automation is that it can be as sophisticated—or as simple—as you need it to be. Our solution deploys in days, not months. Its integrations can be configured and use cases built without developer or coding requirements. As a result, our customers realize their first time-to-value in less than a month.

TAG Cyber: What are the major components of the Swimlane platform?

SWIMLANE: Our vision behind Swimlane Turbine—and the strategy that informed the development of its major components—is rooted in understanding customer needs. We've really noticed four themes emerge. Historically, SOAR technologies were not evaluated on their processing power, but that's about to change, as customers demand more in terms of telemetry ingestion. To address this, we've built an Active Sensing Fabric into Turbine. This technology leverages features like big data ingestion, webhooks, remote agents, pre-processing and inline enrichment to help customers expand their visibility and actionability to the edge. It's designed to ingest data at cloud scale and execute on thousands of concurrent automations in a way that's informed by business logic. Next, customers increasingly need to integrate with things that aren't typically integrated within the SOC, like cloud

Automation can solve any security problem, but it will never reach its full potential if automation continues to be as rigid as legacy SOAR platforms make it.

security, IoT devices and edge-computing infrastructure. Turbine's Autonomous Integrations enable customers to connect to any API without assistance. This environment agnostic approach helps our customers overcome the limitations of closed XDR ecosystems or vendor lock-ins that comes from bundled SOAR offerings with larger enterprise license agreements.

I believe that automation can solve any security problem, but it will never reach its full potential if automation continues to be as rigid as legacy SOAR platforms make it. That's why we are re-thinking security automation, with ease of use and approachability at the forefront. Playbooks are at the heart of this, and Turbine's Adaptable Playbooks pairs low-code condition logic with a human-readable interface to streamline the process of building security automation. Finally, it seems as though security operations is the only business function that does not have a system of record today. Sales has SalesForce, HR has Workday, Marketing Ops has Marketo, but you can't just say security has SIEM. Security needs a centralized platform that combines human and machine data into a composable user experience. Our vision of becoming this type of system of record makes our approach unique. From a platform perspective, Turbine's case management, dashboards and reporting bring this concept to life and help our customers unify their workflows, telemetry and teams.

TAG Cyber: How does your solution address skill shortages and reduce security silos in the enterprise?

SWIMLANE: We help customers **overcome the skills shortage** by making automation more approachable. Turbine's low-code approach enables domain experts to automate without a dependency on developers. This helps analysts automate manual and time-consuming processes so that they can regain time for more strategic work and progress in their careers. Our goal is to enable the average Joe in security to build playbooks faster, so that security teams can adopt more use cases in order to speed their MTTR and increase ROI. Because Turbine is capable of connecting with any API, this unifies disparate technologies and processes in order to **reduce complexity and silos.** This capability means that the platform is truly environment and workflow agnostic, so any organization can connect the dots between their unique workflows, telemetry sources and teams without compromise.

TAG Cyber: Tell us more about how Swimlane handles incident response, threat intelligence, SIEM alerts and other security tasks.

SWIMLANE: While incident response, threat intelligence, SIEM alert triage and other common SOC use cases continue to be where most of our customers start, they are just the tip of the spear. In fact, the inspiration for starting Swimlane actually came out of challenges related to automating a DLP use case back in

my days as a SOC analyst. When we started building Swimlane, we made the deliberate choice to develop the most extensible security automation engine possible. In contrast, most early SOAR vendors had their eye on building specific, yet rigid, predefined SOC use cases. Fast forward to today, Swimlane is now the largest and fastest-growing pure-play security automation company. With the flexibility and adaptability of our platform, we are able to solve these common SOC challenges better than anyone else. We do it by providing a stable, portable and reliable connection to any system in a customer's environment. After this instant connection is established, Swimlane Turbine enables our customer to layer on sophisticated business logic through adaptable playbooks.

TAG Cyber: Can you share some insights into the future of cyberthreats in the upcoming years?

SWIMLANE: Cyberthreats are not going to diminish in the future. Looking ahead, I have three predictions for how cyberthreats will unfold in the coming years. First, **Ransomware will continue to be one of the most common forms of cyberattack.** This is due to the ROI it provides attackers. A recent survey by the Unit 42 security consulting group showed that ransomware payments increased by 82% from 2020 to 2021. Small and medium-sized businesses are especially vulnerable to ransomware attacks, due to the relative immaturity of their security processes. Additionally, **third-party and supply chain attacks will continue to evolve,** especially against government agencies. The U.S. government has made great strides in reducing direct attack threats through various initiatives, along with several guidelines on strengthening responsibility and community-driven threat detection. While the government has revolutionized how they operate from a cybersecurity perspective, federal agencies now have to worry about third parties, instead of direct threats to their services by attackers. Finally, **insider threats will become more common and costly** with an increase of threats motivated by passion or beliefs. Disgruntled employees may be motivated to leak or steal confidential data and intellectual property for their own financial or career gain. The Verizon 2020 Data Breach Investigations Report determined that insiders are responsible for around 22% of security incidents. As these examples show, the security teams of the future will face a higher frequency of common threats, as well as sophisticated new threats that require response processes that go beyond typical SOC workflows.

AN INTERVIEW WITH LORIS DEGIOANNI, CTO AND FOUNDER, SYSDIG SECURITY SOLUTIONS

SYSDIG SECURITY SOLUTIONS FOR CONTAINERS, KUBERNETES AND CLOUD

The transition of application hosting from traditional data centers to modern cloud-hosted infrastructure is rapid and ongoing. Teams now use containers and cloud infrastructure, along with orchestration tools such as Kubernetes, as the basis for delivering application support to their employees and customers.

Cybersecurity company, Sysdig, offers an effective suite of tools that helps application developers, hosting teams and security experts to properly protect modern applications from cyber threats. We had the opportunity to speak with the Sysdig management team recently and were pleased to learn about how their platform works.

TAG Cyber: What types of security requirements exist for containers and Kubernetes?

SYSDIG: The core requirements are similar to traditional security environments, including detection and mitigation for vulnerabilities, misconfigurations, excessive permissions and runtime threats. As organizations build for cloud environments and adopt cloud-native design patterns, security techniques must also adapt, and supporting security tooling must be purpose-built. Modern applications and infrastructure are deployed as micro services, delivered in containers using CI/CD pipelines that accelerate the speed of release and frequency of change. According to the GitLab 2022 Global DevSecOps survey, 70% of DevOps teams release code continuously, once a day or every few days—up 11% from 2021. Developers take advantage of elastic, auto-scaling cloud capacity to build rapidly, but security teams struggle to maintain visibility. Containers are often short-lived, with nearly half existing for less than five minutes, based on Sysdig research. Keeping track of all of an organization's services and configurations is challenging. Ephemeral containers complicate SecOps activities like forensics and incident response.

Compliance and regulatory requirements are also in play, but they must be mapped to traditional infrastructure principles. As an example, the Payment Card Industry Data Security Standard (PCI DSS) defines requirements for cardholder data environment segregation, and organizations must determine how DSS maps to K8s clusters and micro-segmented workloads. Attacker techniques and runtime threats are also unique in cloud-native environments, including cryptojacking and the risk of malicious

As organizations build for cloud environments and adopt cloud-native design patterns, security techniques must also adapt, and supporting security tooling must be purpose-built.

dependencies within container images. These are often described under the umbrella of digital supply chain risks, and organizations are increasingly concerned about the software they consume and use to operate their businesses.

TAG Cyber: What are the major components of the Sysdig platform?

SYSDIG: With our platform, organizations save time on investigation, detection and remediation through a unified view of security requirements from source to run. Security leaders can be confident they have no blind spots regarding critical security threats and misconfigurations in their environment. DevOps workloads can change very quickly, but system calls don't lie. Sysdig has unique visibility at runtime, based on access to Linux system calls. This gives teams a clear view into highly detailed container and host activity for incident response. Curated rules for threat detection are easy to customize. Open-source Falco—the cloud-native runtime security project—is the de facto Kubernetes threat detection engine. Falco provides a straightforward, consistent rules language to ask specific questions against syscalls, Kubernetes audit logs and cloud logs. Sysdig detects malicious activity, misconfigurations and compliance issues through out-of-the-box rules. The Sysdig Threat Research Team meticulously maintains these rules and adds new rules as new attacker techniques are discovered.

Additionally, machine-learning algorithms provide high-fidelity detections. Another layer of detection is enabled by machine-learning models tuned for specific use cases, such as cryptojacking. Framework mapping simplifies threat investigations. Events are mapped to common security frameworks, like MITRE ATT&CK and SOC2, for quick triage. All activity within any application or service by any user across the cloud account, containers and hosts is presented with the detail needed to quickly understand exactly what happened. Runtime intelligence is a game changer for vulnerability management. Risk Spotlight determines if any vulnerabilities found with image scanning are impacting packages loaded at runtime. Prioritization also considers if an exploit is available. By filtering for these factors, the vulnerabilities developers need to address can be reduced immediately by up to 95%. Security policies to validate posture, regarding compliance requirements, benchmarks and internal security mandates are available in our platform. Embracing a Policy as Code approach based on a shared policy model makes it possible to define and enforce security requirements consistently at all stages of the application lifecycle. Policies can be applied through the pipeline, creating a unified toolset, while avoiding tool sprawl and duplicated tasks. The policy model can be used to align security with business and compliance requirements.

TAG Cyber: How does your solution integrate with the DevOps lifecycle?

SYSDIG: Sysdig provides security controls during every phase of the DevOps lifecycle, integrating into the tools DevOps teams use to keep releases moving through the pipeline. Our platform helps teams find and prioritize software vulnerabilities, detect and respond to threats, and manage cloud configurations, permissions and compliance. We provide security for both containers and host instances, including Infrastructure as Code (IaC) scanning at the Git source repository, where teams can see if their IaC configuration is following best practices and compliance updates. We also provide registry scanning and runtime vulnerability scanning of production workloads. Our inline vulnerability scanner permits DevOps teams early access to scan results directly into their CI/CD pipeline tool, which speeds up fixes and eliminates wasted time. Teams can apply a patch or use an updated version of the vulnerable package while building the image, allowing it to ship to production faster.

When Sysdig raises an alert regarding abnormal/risky workload behavior, it can be forwarded to popular notification channels like Slack, MS Teams, PagerDuty and many others, thereby reaching the appropriate DevOps team. Having clarity on the owner of each micro-service application reduces mean time to recovery (MTTR), since teams can start working on the issue right away without the burden of identifying who is in charge. Sysdig also enables DevOps teams to remediate security issues faster with an actionable checklist that aggregates findings to save time during investigations, such as cloud resources that share the host network in an Azure Kubernetes Service (AKS). With Remediation Guru, a suggested fix is provided, improving the CIS Azure Kubernetes Service (AKS) Benchmark score by 200% by just opening a pull request (PR) directly into the source-code repository with the curated patch generated by Sysdig.

TAG Cyber: Tell us more about how your platform works in a modern multicloud environment.

SYSDIG: Sysdig Secure finds and fixes the most critical cloud and container security security risks with no wasted time. This allows DevOps and security teams to efficiently manage risk as they scale their cloud deployment, which is critical given the talent shortage. We provide value through our accuracy in threat detection when using Sysdig's Managed Policies, which is based on Falco and curated by the Sysdig Threat Research Team. We also improve the signals that go into the SOC, which speeds detection and audit workflows in containers, because they get a detailed record of security events. Sysdig can be used across multiple clouds, as well as on-premises Linux environments, as many customers operate in a hybrid mode.

AN INTERVIEW WITH PAUL CIESIELSKI, CRO, TRUEFORT

IMPLEMENTING ZERO TRUST SEGMENTATION WITH TRUEFORT

Application security teams continue to struggle with the challenge of an increased blast radius around their servers. Even as monolithic software in the data center is transferred to containerized designs that are orchestrated using tools such as Kubernetes, the protection challenge continues, especially if code changes are required to enforce security.

TrueFort offers a novel Zero Trust segmentation approach that secures applications in the cloud and data center without requiring code change. We asked the TrueFort team to help us understand their platform solution, including how it works with best-in-class endpoint agents.

TAG Cyber: What are the greatest risks to cloud workloads?

TRUEFORT: While most workload-protection vendors would claim that the biggest risk to every organization is the compromise of a single cloud workload, the reality is that it has little impact on an organization. The greatest risks to cloud workloads revolve around the tactical advantage they offer an attacker with access to one. Once reached, attackers are often able to operate for days or weeks without detection, because they are already inside the perimeter, and the understanding of normal operations is especially weak. Moreover, on most workloads, the execution of important commands is automated with privileged accounts, meaning attackers can steal these accounts from the compromised systems and reuse them to move laterally. Finally, from an initial cloud workload access, attackers easily move from one asset to another in order to access critical data stored in a database within the same environment.

TAG Cyber: What specifically is meant by Zero Trust segmentation?

TRUEFORT: Historically, all access to a production environment required authentication, but after a user was authenticated once via VPN, they were trusted to access any workload or database within the same environment, without limitation. Attackers famously took advantage of this by stealing privileged credentials and using them to rapidly spread across the network, system by system, until they reached a database containing critical data they could monetize. Attacks may originate as phishing successes or SQL injections on a website, but they escalate when they gain administrator control of a cloud workload.

Microsegmentation is rapidly moving from a "nice to have" to a "must have," when it comes to protecting both critical and regulated data.

Zero Trust aims to disarm this common attacker behavior by questioning every single lateral move within the data center or cloud. The two major aspects of a Zero Trust architecture are: Zero Trust Network Access (ZTNA) and Zero Trust segmentation. ZTNA is determined to improve upon VPN access by testing every access of an administrator within the data center or cloud, whereas Zero Trust segmentation focuses on making it impossible to move from one workload to another in a manner that is not necessary for typical enterprise applications to serve customers. More specifically, Zero Trust segmentation is an aspect of Zero Trust that blocks access from one workload to another between, and within, enterprise applications. This means that despite an account getting approved to access one workload in the environment, all access to a second or third workload is blocked until reevaluated and approved, thus creating a significant number of network segments within a single application environment.

TAG Cyber: How does the TrueFort solution work?

TRUEFORT: Our solution makes Zero Trust segmentation (or microsegmentation) easier and more effective by starting with a clear behavioral mapping of all activity within the data center and cloud from Day One. Rather than merely demonstrating network traffic between workloads, TrueFort shows security teams all activity according to the **who** *(Which service or admin account?),* **what** *(What command was executed?),* **when** *(Does this happen often?),* and **where** *(Is the resulting activity at the destination unusual?).* Once our customers have this clarity, they are able to not only enforce the blocking of network connections through host firewall rules, but also automatically kill any of these behaviors that are unapproved. They can prevent processes from running, shut down an account behaving strangely, or even go so far as to kill an unusual command line argument to instantiate a network connection that's been witnessed thousands of times before.

TAG Cyber: Tell us more about how you integrate with commercial endpoint agents.

TRUEFORT: In general, one of the biggest pains around security products is that they always require "yet another agent." This is why we invested so much development effort into making the TrueFort Platform work with existing CrowdStrike and other endpoint detection and response (EDR) agents that our customers have already deployed. The full extent of how service accounts are used to execute key network connections between applications and workloads is gleaned through the telemetry gathered by agents that are already installed. And it doesn't stop with analyzing the telemetry that comes from these EDR leaders, TrueFort also pushes enforcement policies to the agents

to automate microsegmentation. As IP addresses change from the DHCP or new workloads spin up to auto-scale for a customer, TrueFort fingerprints these workloads based on their behavior within enterprise applications and uses that to identify anomalous actions. No one in the cyber security market shows more value from the EDR agents that customers have already installed.

TAG Cyber: What predictions can you share regarding this area of enterprise cybersecurity?

TRUEFORT: From what enterprises are experiencing, it is clear that microsegmentation is rapidly moving from a "nice to have" to a "must have," when it comes to protecting both critical and regulated data. Since it is a likely occurrence that there will be some level of compromise in any application environment, microsegmentation is now viewed as the only viable mitigation against the lateral movement that makes these attacks so devastating. Between NIST and the DOD both including microsegmentation as mandatory for Zero Trust architectures, cyberinsurer mandates, and regulators requiring deeper segmentation for compliance, it is time for every organization to build their plan for enhancing enterprise security with microsegmentation.

"My Mom does cybersecurity and I seriously think she has trust issues."

AN INTERVIEW WITH TOM GILLIS, SENIOR VICE-PRESIDENT/GENERAL MANAGER, NETWORK & ADVANCED SECURITY BUSINESS GROUP, VMWARE

MODERN VIRTUAL AND CLOUD SECURITY SOLUTIONS FROM VMWARE

VMware has been a major pioneer in the delivery of virtualized applications, systems and infrastructure to modern enterprise customers. With this innovation, the need has arisen to integrate a suite of powerful security solutions that can work with VMware installations, as well as any adjacent systems deployed into or around the VMware software.

More specifically, VMware offers built-in and distributed security solutions that control threats to users, devices, workloads and entire networks. The goal is to minimize risk, while also reducing the need for corporate silos. We enjoyed spending time with the VMware team to gain valuable insights into how their platform is evolving.

TAG Cyber: Can you give us an overview of the more general VMware solutions offered to enterprise customers?

VMWARE: As you already know, the flood of cyberattacks has continued into 2022, with persistent and increasing reports of ransomware. A big part of the problem in mitigating these attacks is that attackers have learned to gain access using legitimate pathways and then move laterally within an enterprise's infrastructure easily, as there are little or no barriers restricting their movement. The good thing is, VMware solutions addresses this by providing "lateral security." Through the power of virtualization and software, we apply best practices of well-known security capabilities, making them more pervasive throughout the infrastructure. For example, VMware's flagship product, NSX Firewall, addresses this through the compartmentalization of each network segment. NSX Firewall enables enterprises to detect and prevent lateral movement within permitted traffic. In addition, NSX Firewall is backed by threat intelligence from VMware Contexa, allowing it to keep up with the latest attack methods.

TAG Cyber: What are the major components of the security aspects of the VMware platform for customers?

VMWARE: When we talk about NSX Firewall, there are two components in play here: NSX Distributed Firewall and NSX Gateway Firewall. The first, NSX Distributed Firewall, is a distributed, hypervisor-resident internal firewall that protects

Attackers will continue their devastating attacks until enterprises erect significant barriers with "lateral security" solutions.

each workload. Enterprises use it for network segmentation and microsegmentation to reduce permitted network traffic. NSX Distributed Firewall comes with Advanced Threat Prevention (ATP) that provides intrusion detection/prevention systems (IDS/IPS) and network sandboxing, as well as network traffic analysis (NTA) and network detection and response (NDR) services for a complete network security stack. Obviously, it is ATP that provides a line of defense against threats inside permitted traffic. The second component, NSX Gateway Firewall, is a software-based firewall typically deployed at the perimeter of an environment protected by NSX Distributed Firewall. All together, the Distributed Firewall and Gateway Firewall secure the entirety of private and multicloud environments, while presenting a consistent management interface to networking and security administrators.

TAG Cyber: How does your solution integrate with other cloud infrastructure and existing, deployed security tools?

VMWARE: I'm glad you asked because a primary design goal of the NSX Firewall is to support the cloud operating model. That is, NSX Firewall aims to provide the public cloud's speed and agility in secure private and multicloud environments. This is a solution that can be managed entirely via APIs. In fact, anything that can be done with NSX's management console can also be done with the API and vice-versa. These APIs enable enterprises to deploy secure workloads with zero IT tickets. Further, NSX Firewall provides extensive logging in a standardized format. These logs are ingested by VMware and third-party tools already deployed by enterprises. Finally, NSX Network Virtualization—a sister product of NSX Firewall—enables enterprises to insert third-party security tools in the path of network traffic. This means that enterprises with these actively deployed third-party tools can continue using them alongside NSX Firewall.

TAG Cyber: Tell us more about how your full range of customers—from large enterprises to small and mid-sized companies—benefit from partnering with VMware on security.

VMWARE: One of the best benefits of VMware is our enormous penetration of small, mid-sized, and large enterprises with server virtualization. We also have a growing presence with enterprises that containerize their workloads. Thanks to all of this, VMware's NSX Firewall is attractive to all of these enterprises, irrespective of size. The first reason is that NSX Firewall is a mature offering with commercial availability for virtualized workloads dating back to 2013. The second is that VMware was an early adopter of the trend toward containerization. As a result, NSX Firewall has robust native support for containerized workloads.

Most enterprises are heavily invested in virtualized and containerized workloads. They have also acquired point security products to cobble together a homemade lateral security solution for their workloads. However, most of these homemade solutions are incomplete in terms of depth of deployment or completeness of network security stack, or both. Enterprises that invest in understanding and deploying the NSX Firewall can transform their security posture by creating significant barriers to lateral movement. As a result, the effort and cost of partnering with VMware to deploy the NSX Firewall are dwarfed by the potential effort and cost of recovering from a cyberattack.

TAG Cyber: Can you share some insights into the future of cyberthreats in the upcoming years?

VMWARE: Absolutely. For starters, enterprises are well defended at the perimeter with perimeter firewalls. Many enterprises have also acquired modern endpoint security solutions and installed identity and access management systems. At the same time, attackers have become adept at probing for and exploiting weak points in an enterprise's defenses. Invariably they find a way to get a foothold on an enterprise system. The real problem lies within: Most enterprises present only weak security barriers to attackers inside their infrastructure. As a result, attackers can rapidly expand from their initial foothold to a pervasive and persistent presence inside the attacked enterprise's infrastructure. Eventually, these attackers inflict enormous damage to enterprises via ransomware or leaked intellectual property. Attackers will continue their devastating attacks until enterprises erect significant barriers with "lateral security" solutions. This is VMware's main point. Of course, VMware aims to provide enterprises with the best lateral security suite to erect these barriers.

"Next Generation Access Control."

AN INTERVIEW WITH RAVI SRINIVASAN, CEO, VOTIRO

A ZERO TRUST CONTENT SECURITY SOLUTION FROM VOTIRO

Enterprise security teams need a cybersecurity solution for safely handling unstructured data files, including the need to remove any hidden malware (unknown, evasive) quickly and at scale. Instead of relying on detection technologies, a proven solution exists called content disarm and reconstruction (CDR). The CDR process eliminates malware found in files without using detection. For all modern digital and cloud-based services, CDR is best done as a SaaS capability to ensure rapid use and scaled operation.

Cybersecurity vendor Votiro offers an effective CDR solution that is delivered through an open API-based SaaS, consistent with a new concept known as Zero Trust Content Security. We wanted to learn more from Votiro about the evolution of CDR, and how Zero Trust Content Security can benefit an enterprise.

TAG Cyber: What are the most effective modern methods for removing malware from files, content and data?

VOTIRO: Weaponized files are the easiest way to infect an organization, and they can do so undetected. In a recent report from "Cybersecurity Insiders," nearly 70% of malware found in files is of unknown variant. That means file-based attacks using techniques such as phishing, zero-day threats, and ransomware are still the easiest form of successful attacks. Think of how you as a *USER* interact with files (PDFs, DOCs, PPTs, ZIPs) daily. You use email, web browsers and collaboration platforms. Now, think of all the ways *APPS* automatically exchange files with supply chain partners and/or receive files from any source in the cloud. All of these channels that ingest files also widen the attack surface and open the risk of file-borne threats entering. There are many detection-based solutions that attempt to tackle file-borne malware, but simply looking for known malware signatures using multiple AV engines, anti-malware tools or sandboxes is not effective and leaves protection gaps. There is a time lag between a new malware variant being created—or even just a simple tweak to an existing one—and the signature being recorded in databases that detection-based solutions pull from. Within that gap of hours, days or weeks, there is the space for that malware to enter an organization, move laterally, exfiltrate data, deploy ransomware, or wreak other havoc.

Therefore, leading global governmental and commercial organizations are adopting a proactive approach to preventing evasive and unknown malware: **Content Disarm and Reconstruction (CDR)**—a Zero Trust content

Instead of looking for malware and blocking it, CDR cleanses and delivers safe files to users and applications.

security process that deconstructs *every* file, cleanses it and delivers the now-safe file to users and applications—wherever the file was originally headed. What makes CDR effective is that users and applications get safe files in real time to use in their respective business processes: email, web downloads, secure file transfers, collaboration and cloud storage. After a file is cleansed using CDR, all the left-over bits such as unknown objects, embedded code and potentially malicious executables can be proactively analyzed by security teams, eliminating the need for costly detection and response cycles.

TAG Cyber: What are the major components of the Votiro platform?

VOTIRO: Votiro Cloud offers open API-based, cloud-native services that enable security and digital teams to analyze 150+ ISO-standard file types, including audio/video files in audit mode and deploy CDR in public SaaS or private cloud mode, closest to where the data is used. Votiro Cloud seamlessly integrates with existing security, application and data management platforms. Most common integrations include deploying CDR as a service with business email compromise, secure file transfer, remote browser isolation, cloud storage, collaboration and data management platforms.

TAG Cyber: How does your solution ensure that enterprise security teams are never slowed down and that content, data and file flows are preserved?

VOTIRO: This is the primary value proposition for CDR: Instead of looking for malware and blocking it, CDR cleanses and delivers safe files to users and applications. **Our CDR delivers fully functional, usable and safe files.** Votiro applies a patented Positive Selection Technology—a process that starts with a clean, ISO-standard template of the target document format, then populates all usable content into the clean template, tests for file fidelity and delivers the fully functional file to the target destination. This process happens in **milliseconds**. Votiro is unique in that we don't flatten documents or strip active content from the documents, including benign macros used in business documents. Our CDR implements advanced machine-learning techniques to include safe macros in office documents and safe images with built-in anti-steganography capabilities. Votiro helps enterprise security operations reduce the burden on their detection and response efforts by implementing a RetroScan capability to analyze potential threats after safe files are delivered to the destination, as well as offer native integration with common security information and event management (SIEM) platforms.

TAG Cyber: Tell us more about the Votiro Cloud and how it works.

VOTIRO: Votiro Cloud is an open, API-based service that can be easily integrated into an enterprise's existing security and application environment. Here's how it works: Votiro CDR is a run-time, proxy service that supports email API, browser plug-in, and application API, providing ease of integration, while eliminating the need for additional security controls on the endpoint. Votiro also supports flexible deployment modes—public SaaS and private cloud that meet enterprise IT and compliance requirements.

TAG Cyber: Can you share some insights into the future of cyberthreats in the upcoming years?

VOTIRO: What has the past 20 years of cybersecurity experiences taught us? **Content is king**—and this is still very true for today's modern digital businesses. Good guys need to use it to be productive and create value. Bad guys need it to cause harm and business disruption. I see the future of cybersecurity as applying Zero Trust principles to protect the CONTENT that matters. Specifically, Zero Trust principles will need to go beyond CONNECTIVITY AND ACCESS to securing the use of DATA. Zero Trust Content Security will allow digital services to use CONTENT safely. Government and commercial organizations will focus on safely using CONTENT in their everyday business processes. CISOs will have to focus on the CONTENT entering the organization's application, data management and business services that are not visible to the infrastructure-centric security. For example: How do you safeguard CONTENT while: exchanging files with partners using public cloud storage like Box or Drive; uploading files into data lakes like AWS S3 buckets; downloading files in public channels within collaboration tools like Slack; or transferring files via APIs into applications and data management platforms? I see the promise of implementing Zero Trust Content Security as a proactive and preventive measure to reduce the risk, cost and burden on security teams.

ANALYST REPORTS

ANALYST REPORT

USING CYBER RISK INTELLIGENCE TO MANAGE THIRD-PARTY CYBER RISK: AN OVERVIEW OF CYBERGRX

DR. EDWARD AMOROSO

Cybersecurity risks of third parties can be addressed effectively using cyber risk intelligence management. The CyberGRX Exchange supports third-party cyber risk management and accelerates the assessment completion and evaluation process for enterprise teams and their supplier partners.

INTRODUCTION

Security practitioners agree that third-party cyber risk management (TPCRM) has emerged as a **primary concern** for chief information security officers and their teams. Justification for the concern is easy to establish: whereas security teams can exert control over the protection practices in their own enterprise, extending any type of security control to suppliers, partners and even customers can be considerably more difficult.

The traditional approach to TPCRM relies on the sole use of questionnaires—albeit only for larger suppliers. (Many enterprise teams levy no security requirements on their smaller third parties.) While the idea of a questionnaire is familiar, the process requires considerable time and effort from third parties, especially when they must answer the same questions repeatedly.

In this report, we describe the modern challenges that exist for enterprise teams dealing with TPCRM, as well as measures that have been used to address the problem. We then focus on how a data intelligence approach can be used to manage cyber risk and the reputation of third parties. The commercially available **CyberGRX Exchange** is used to show how such a solution can work in practice.

CHALLENGES OF THIRD-PARTY SECURITY

Most enterprise security teams will agree that while it is obviously challenging to manage *their own* cyberprotection posture, it is even *more challenging* to extend such posture assessment to third-party environments outside their control. This is especially true for larger third parties which include the usual cyber risk complexities, and which also must address the challenge of dealing with their own third parties.

From the perspective of the third-party company, similar challenges exist. That is, the typical supplier or partner working with a buying entity can find that the security requirements being imposed are onerous. Third parties are rarely sufficiently resourced to handle the demand for multiple security assessments, and they must often contend with unreasonable challenges to respond quickly with detailed information.

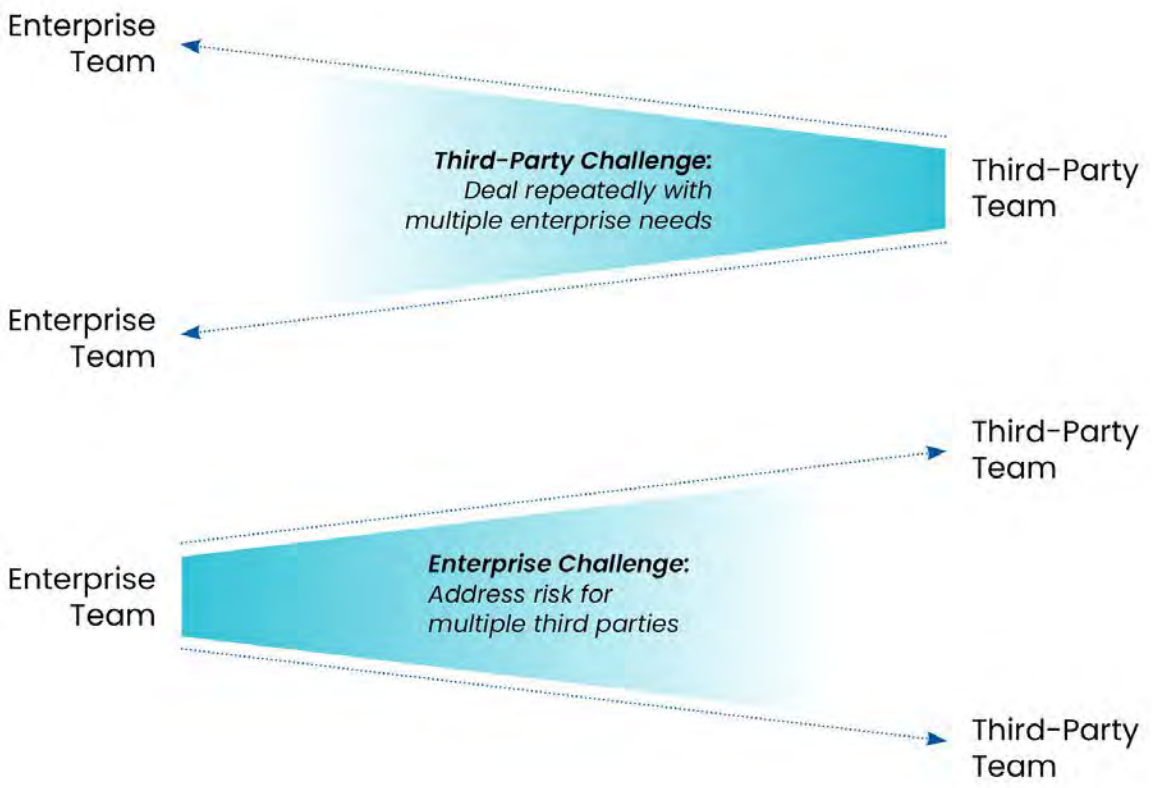

Figure 1. Challenges of Third-Party Security

Because the challenges here are significant, new approaches are emerging that ease the difficulties for both enterprise teams dealing with multiple third-party risk and suppliers and partners who must deal with the security requirements being levied by multiple enterprise teams. In the next section, we outline the basics of an approach that is based on the use of *data intelligence*.

ADVANTAGES OF A DATA INTELLIGENCE APPROACH

A new approach to third-party risk management uses data intelligence for visibility into cyber risk and reputational profiles. The goal is to use data analytics to accurately assess third-party risks so that your third-party portfolio can show you a more holistic understanding of the security posture of those vendors. This view can ease procurement and improve the TPCRM process. The key data intelligence features for TPCRM include the following:

- ***Emphasis on Data Analytics***—Rather than just managing questionnaires, risk-intelligence-driven approaches depend on the use of advanced analytics. This process works best when the data collected forms a repository that can be used to validate existing practices and predict future issues.
- ***Support for Scaling***—For TPCRM to scale across enterprise and third-party participants, methods such as crowdsourcing can simplify data gathering. An additional strategy is to create an exchange which can reduce the amount of coordination required between multiple enterprise and third-party companies.
- ***Focus on Actionable Insight***—Effective solutions for TPCRM must focus on driving actionable insights for both enterprise and third-party participants. This implies that the repository and analytics support tasks such as prioritization, benchmarking and other methods of assessing or reducing risk.

These high-level requirements should be included by enterprise security teams in their solicitations for third-party risk management support. They are designed to complement the well-known scoring and questionnaire methods that are so commonly included in commercial platforms supporting this work. In the next section, we provide an overview of CyberGRX, which effectively addresses these requirements.

CYBERGRX PLATFORM OVERVIEW

Founded in 2015 and headquartered in Denver, the CyberGRX Exchange is revolutionizing how security practitioners think about TPCRM and protect their organizations. Enterprise organizations rely on CyberGRX to simplify procurement, support regulatory demands and improve their cyber risk management process. The specifics of the CyberGRX Exchange are explained below:

Cyber GRX Exchange
The Exchange is the foundation for the platform offering. It collects assessment data from third parties in a structured format to create data that assists in analysis, reporting and comparison of third parties. This data provides actionable insights for companies at both a portfolio and individual third-party level. Ultimately, the security goal is to learn where third parties are lacking and to either change suppliers or help them mitigate any relevant cyber risk.

Framework Mapper
This component of the platform supports mapping of the CyberGRX assessment to industry frameworks. Such analysis enables more granular understanding of where a given third-party supplier might have its security strengths and weaknesses. Popular security and privacy frameworks covered by the CyberGRX Exchange include NERC, CMMC, NIST 800/CSF, HIPAA, PCI-DSS, CCPA, GDPR and SIG.

Threat Profiles

This component is focused on the security controls that are implemented in a third-party enterprise. By understanding these controls, teams gain insight into how and whether a given supplier might be affected by common attack methods. It complements the Framework Mapper through analysis that is focused on recent attacks that have been seen in live practical environments. Threat profiles are built based on the MITRE ATT&CK framework.

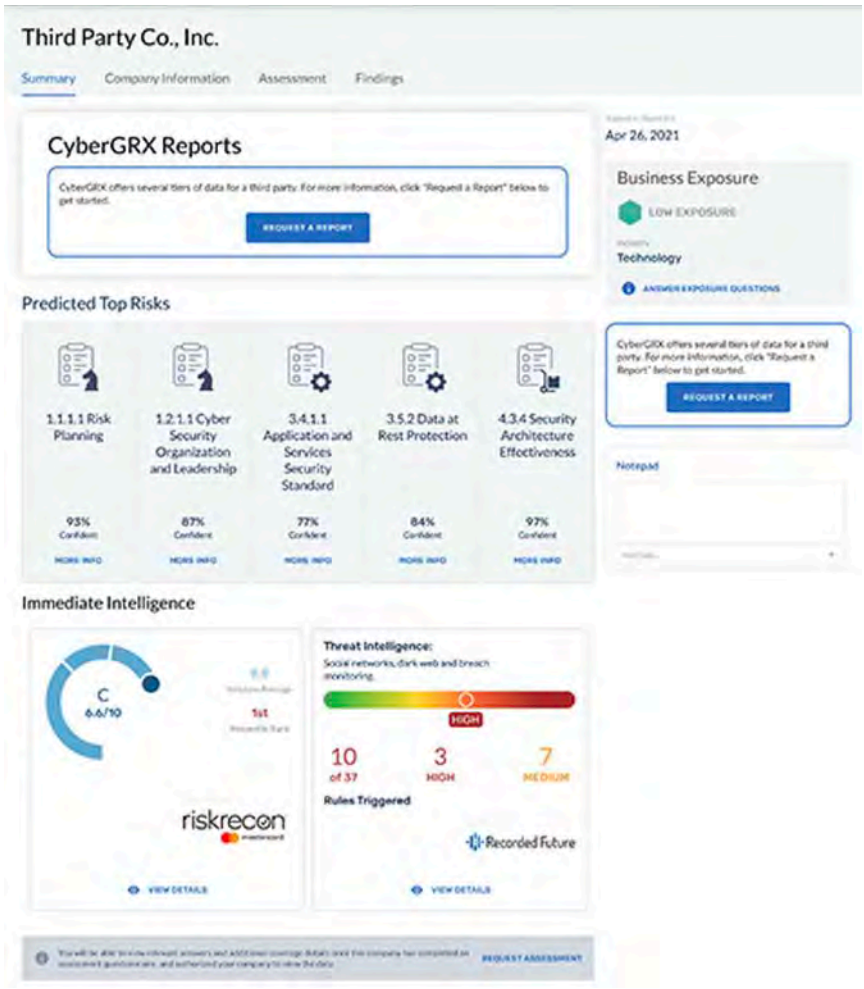

Figure 2. Sample CyberGRX Predictive Risk Profile

Predictive Analytics

CyberGRX applies machine learning algorithms to its data to generate predictive models about third-party cyber risk. This is a powerful means to enable predictive processing. In fact, the technology generates a future view of how a supplier will answer questions based on factors such as similar completed assessments. This is the right way to use machine learning to reduce risk, and CyberGRX reports up to a 91% accuracy rate on these profiles.

Using these important platform features, CyberGRX generates a *Predictive Risk Profile* (see Figure 2) that can be viewed as a measure of cyber risk and reputation. This is obviously important to enterprise teams engaging third parties, but it is also a powerful metric for companies to drive improvements in their own infrastructure. Having a strong reputational profile is good for business and should lead to more contracts and higher revenue.

TAG CYBER

ANALYST REPORT

MODERNIZING THE SECURITY INFRASTRUCTURE: MITIGATING ENTERPRISE CLOUD WORKLOAD RISKS IN LEGACY INFRASTRUCTURES

JOHN J. MASSERINI

Cyber security risks associated with using workloads on Amazon Web Services, Google Cloud Platform, Oracle Cloud and Microsoft Azure are distinct and inherently different from legacy enterprise infrastructure risks. Unfortunately, most security teams are ill-equipped to fully understand the risks of these cloud environments and how posture management, vulnerability mitigation, user access and data management are fundamentally different from the typical enterprise paradigm.

In this report, we will review the challenges of leveraging a modern cloud-based workload infrastructure, how it differs from its legacy peer and how the Sonrai platform can highlight these new risks.

INTRODUCTION

Moving to cloud-based, virtual workloads is a primary driver in many of today's enterprises. One **recent study** by the MIT Sloan School of Management reflected how public companies who are aggressively adopting cloud strategies have a 2.3%-6.9% higher annualized revenue growth rate. **Yet another study** has shown that almost half of the enterprise compute work effort is performed within the Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure (Azure), or Oracle cloud environments.

While the benefits of a cloud strategy to enterprises are substantial, the risks which accompany these modernization efforts are often misunderstood or misrepresented not only within the security teams but across the organization due to shifting paradigms and the related gaps in coverage of the existing security tools. Unlike in legacy infrastructure environments, where security teams could rely on a high degree of visibility from design to implementation as well as some degree of direct control, cloud infrastructure migration, along with the move to DevOps, empowers the development teams to quickly stand-up entire environments, which include workloads, data stores of various type, as well as networks and cloud-based firewalls, with zero reliance on the existing change management processes. Gone are the days of architecture reviews and security evaluations for new applications and requesting firewall changes, DNS updates, or server configuration changes – today DevOps teams can deploy a fully integrated application suite in mere hours, without any assistance/approvals from, and often the knowledge of, the security team.

In this report, we will analyze the stark differences in risks of today's cloud environments in comparison to the legacy environments relied upon by most enterprises over the past 3 decades or more. Throughout the report, we will be highlighting the critical gaps in coverage typically seen by enterprise security teams and ultimately, how the Sonrai platform can facilitate the management of said risks.

ADVANTAGES OF THE NEW MODEL

The advantages of the cloud-based model are numerous. As an example, if we look at AWS, the ability to stand up a 'server' (i.e., workload), allocate a data store to it via an AWS S3 bucket, and assign it to a network with full internet access can now be done by a developer – with zero dependence on any other operations or security teams. Gone are the days of waiting weeks or months to stand up a test environment for developers. Nowadays, a database can be instantiated, tables loaded with data, and cloud-based firewalls opened to the internet – all within a couple of hours. All of this can be done with code, at a scale and speed that is unheard of outside of the cloud.

Today's rapid application deployment models are not only enabling a transformative business climate but are inherently driving significant technology modernization throughout many enterprises. The result is that the DevOps process has turned not only the infrastructure acquisition process but the entire process of building applications and whole environments on its head.

Not only does leveraging cloud solutions benefit the development process, but the ability to leverage virtual infrastructures rather than dealing with lengthy hardware supply-chain delays is forcing internal IT infrastructure teams to re-evaluate the necessity of having on-premises solutions. 'Cloud First' initiatives have not only become a nice-to-have but have evolved into a mandatory requirement to keep the business running.

While not typically a high-risk priority such as a data breach, cost containment is a factor in every enterprise. While many look at cloud solutions with an eye toward cost savings, the tendency of the environment to sprawl should be strongly considered. Oftentimes, temporary files, databases, and code repositories are left on abandoned workloads which, either running or shut down, consume storage and therefore, the fees associated with it. Having the capability to inventory and audit the cloud environment is no longer a nice-to-have, but an unquestionable necessity.

CHALLENGES OF CLOUD-BASED INFRASTRUCTURE MODELS

While most security executives have accepted the existence of the cloud within their organizations, most fail to understand the substantial difference in risk posture within this new model. In the previous example, while the developer was able to deploy their application stack quickly and easily, neither the

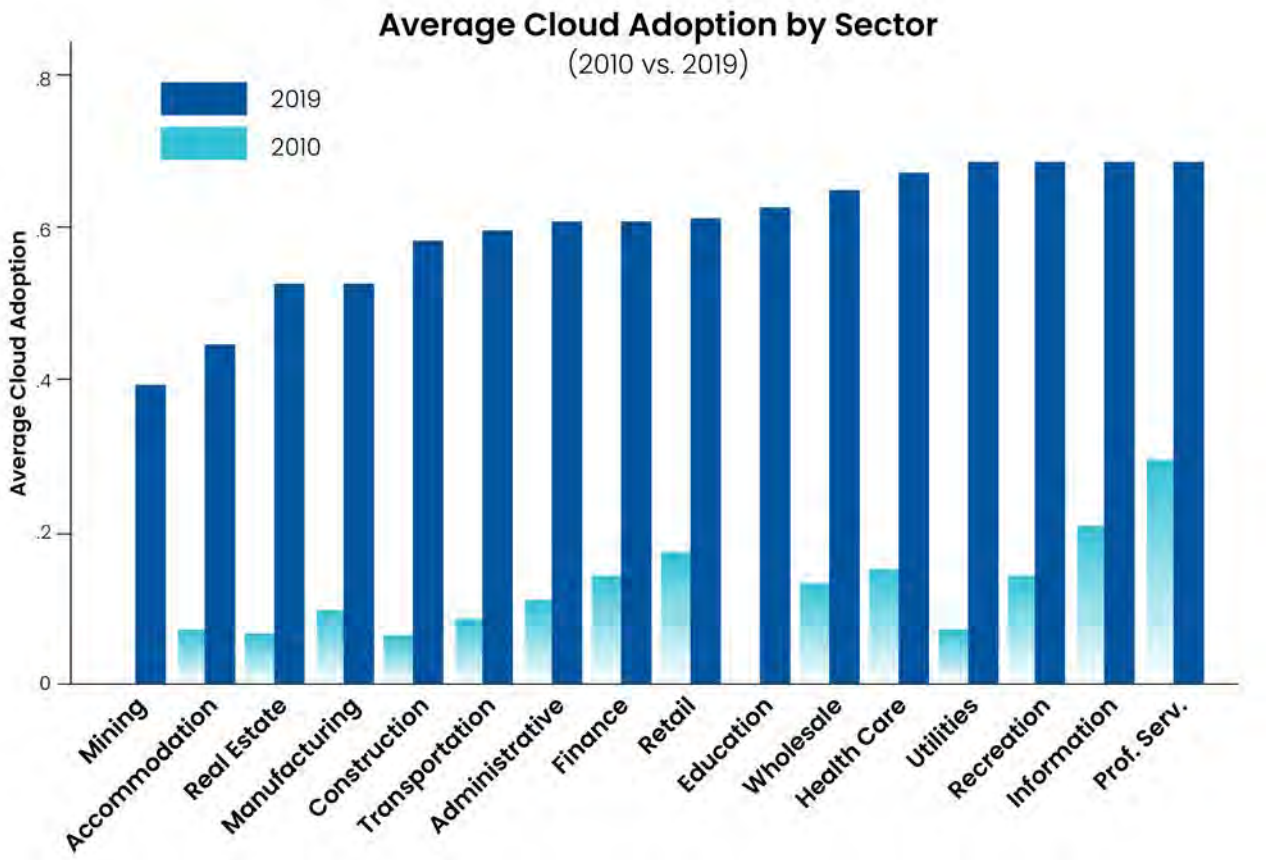

Figure 1: Cloud adoption 2010-2019 @Wang Jin

security nor the infrastructure operations teams were alerted to the existence of the environment, much less given an opportunity to review the posture or the deployment settings of not only the workloads but of the entire environment.

Unfortunately, there are two fundamental root causes of this lack of visibility. The first is the inherent reliance on legacy process management to identify when changes occur within the infrastructure. Consider the average enterprise: when a business wants to offer a new application or service, they go through a rigorous process of architecture design & engineering, and hardware procurement, deployment, testing, and then operationalization. Each step along the way documents the environment changes, identifies those who are responsible, and ensures all ITSM change management efforts have been completed. From IP address assignment to firewall or DNS changes, the infrastructure operations teams 'own' the process and can ensure that new devices meet the documented requirements, have the appropriate controls, and do not disrupt the stability of the environment. This 'catch-all' legacy hardware and application deployment process is tried and true – and fails miserably when cloud technologies are introduced to the environment.

While an accurate device inventory has always been a challenge for most technology departments, the introduction of the cloud, with the ability to spin up workloads at a moment's notice, make an already problematic issue substantially more onerous. While most public cloud providers offer tools to manage inventory, all too often the function is neglected due to the lack of most companies assigning an individual to own cloud usage. Because of this, Asset inventory is virtually non-existent when it comes to most enterprise cloud environments, even though they have the ability to do it to an extent far greater than ever before.

The other major issue most security teams are facing is the inability of most legacy security tools to identify the risk and security issues around cloud workloads and applications. From configuration management, vulnerability scanning, or supplying security logs to the SIEM/SOC, cloud environments are often overlooked by the security teams mainly because of their inability to see what is running within the environments. Additionally, even when legacy security tools are deployed, their lack of conceptualizing the entire cloud environment leaves them woefully inadequate when it comes to identifying cloud-based risks.

TECHNICAL AND DATA MANAGEMENT RISKS OF CLOUD ENVIRONMENTS

Fundamentally, cloud risk management breaks down into two distinct risk profiles, Technical Security and Data Management. The technical security risks of most concern are three of the same issues you would find in a typical infrastructure; misconfigurations, vulnerabilities, and access controls/excessive permissions. As for Data Management risks, the lack of controls over data stores and the information stored in them is of critical importance within cloud environments. Not only of concern is the sensitivity of data stored, but also how it's stored, who or what has access to it - and what permissions have been provided - as well as what they are doing with the data.

TECHNICAL SECURITY RISKS

In a similar vein to their legacy counterparts, cloud environments are vulnerable to the same types of technology-centric risks that security teams have been dealing with for decades. Unfortunately, due to the inherent trust relationships which exist in most cloud environments, risks that would have historically been limited in scope to specific devices or networks now impact the entire platform. While such trust relationships enhance the speed by which deployment can occur, it also greatly increases the possibility that an otherwise innocuous user or vulnerability could be leveraged to subvert the entire cloud environment. Trust relationships within the modern cloud environment are truly their Achilles Heel.

ACCESS CONTROLS

One of the most critical issues facing cloud environments is the lack of controls around user and machine access management. Having common, shared passwords across users and systems is a frequent deployment method to ensure everyone who needs access can get it. In general, the DevOps process consists of a write-build-test-deploy cycle called a 'sprint' which generally lasts around a couple of weeks. To ensure everyone on the DevOps team, from the developer to the testers, to the business partner, has access to the applications as needed, common, easily shared non-person cloud identities, such as AWS Roles, Azure Service Principles, etc., are used. Unfortunately, such access rights tend to migrate into production as the application does, providing access that was intended for testing now with full permissions to often sensitive production data.

This same issue exists when you evaluate the trust relationships in workload-to-workload (machine-to-machine) connections and communications. If we were lucky and the developer decided to configure SSH connectivity or something as basic as FTP accounts to move flat files from host to host, SSH keys and accounts with weak passwords or excessive privileges tend to follow the system all the way to production. As previously highlighted, the lack of ability to track and report on these users and privileged keys leaves the enterprise highly vulnerable to both external and internal attacks.

POSTURE AND VULNERABILITY MANAGEMENT

Much in the same way as access control issues mirror that of the legacy enterprise, so do risks associated with posture and vulnerability management. If a developer decides to deploy a LAMP (Linux-Apache-MySQL-PHP) stack, they will likely choose the image that gives them the greatest number of development

and application options, rather than the one that is pre-hardened. Many of these workload images are designed for ease of deployment and take few security controls into consideration, resulting in a highly vulnerable system that is now being exposed to the public. As we've highlighted previously, the lack of visibility into the workloads, or even that the workloads exist, leaves the security teams behind the proverbial 8-ball when it comes to trying to manage risks in the cloud. This has never been more apparent than after the recent Spring4Shell and Log2Shell vulnerabilities, where without an accurate workload inventory, determining the risk to the enterprise would be nearly impossible.

Unlike in a legacy environment, however, vulnerabilities and configuration issues on workloads can result in a significant risk of breach across the entire cloud environment. Time and time again, vulnerabilities have been used to exploit misconfigured workloads, resulting in the attacker gaining access to a specific workload. While significant, due to the frequent misconfiguration of identities and roles within the cloud environments, the attacker now potentially has elevated access to every other workload within the environment.

DATA MANAGEMENT

While generally, the focus of most risk analysis tends to be on vulnerabilities and threat actors, the way data is managed throughout a cloud environment can substantially mitigate those threats, or conversely, elevate the risk if not managed correctly. Due to the nature of the DevOps lifecycle, the various data stores within cloud environments tend to replicate quickly, leaving outdated databases or file extracts lying around unprotected. A solid cloud data management program is essential to mitigating data breach risks of any cloud deployment.

When developing a data management program for most cloud deployments, the security teams must consider four key elements:

- Where within the cloud is the data stored? (Location)
- What type of data is stored in the cloud environment? (Classification)
- Who /what has access to that data? (Entitlements)
- What are they doing with it? (Usage)

Due to the fundamentally different ways in which most cloud environments are managed, understanding the inherent risk associated with cloud data storage should be the foundation of any cloud strategy.

Production Data Management

Protected data, be it customer information, health care or financial records, or any type of intellectual property, must be managed especially tightly within cloud environments. Due to the lack of segmentation, the hard, air-gapped network boundaries that most infrastructure teams rely on disappear within cloud environments.

Since the major appeal of adopting DevOps is the rapid time-to-market for new business applications, understanding how production data is used and managed is key. Much like the applications themselves, data models change frequently throughout the early DevOps sprints so models which had little protected data in the early stages will often end up with a collection of elements that elevate the entire cloud data store to a protected level.

As more enterprises adopt a DevOps model, managing production data and appropriately mitigating these new types of risks will be critical. This gap was a root cause identified in one of the most

noteworthy breaches of the last several years. As the after-action report from the Department of Justice noted, the **Capital One breach** in July 2019 was due to a vulnerable workload, which had access to cloud data stores of customer data, that was exposed to the internet.

Test Data Management
Within any typical development shop, regardless of the methodology used, the need for accurate, current test data is critical in ensuring the applications are functionally tested adequately. As such, duplicating, moving, and backing up databases is a common occurrence in development environments. Unlike in legacy infrastructures, cloud-centric development teams have the capability of making countless copies of data stores, datasets, full databases, or flat files, full of the customer data

SONRAI PLATFORM OVERVIEW

Sonrai offers a total public cloud security solution for Amazon Web Services, Microsoft Azure, Google Cloud Platform, and Oracle Cloud. Sonrai Dig identifies and monitors all relationships between workloads, identities, and data stores that exist within your various cloud platforms to provide security teams a continuous view of all risks, unusual activity and automated remediations.

Through a single user account, Sonrai can analyze host, datastore, and user activity to provide a complete picture of what's happening across your clouds. Additionally, Sonrai applies analytics to the workload activity logs to provide deep insight into data access activities and user permissions across the entire environment, uncovering risky relationships between workloads and identities.

By leveraging the Sonrai Governance Engine, security teams can automate risk remediation workflows such as patching, posture validation, or user access/privilege management to ensure the platforms stay secure in real-time. By being able to identify protected information within the various datastores, Sonrai can determine the true risk to your data across the entire platform, not just based on the risk of a single workload.

ACTION PLAN

Surprisingly, the 'not in my backyard' opinion of many security leaders is indicative of the cloud security problem, and unfortunately, far too many have ignored the rapid expansion of such platforms within their environments. While most cloud platforms offer adequate security controls, more often than not, developers choose ease and speed over security since many are left to make their own decisions.

The DevOps model, and cloud solutions in general, continue to be adopted at an ever-increasing pace. As business leaders recognize the increased revenue growth associated with cloud strategies, the Information Security teams must find ways to adapt and, more importantly, participate in this rapid deployment approach. Ongoing reliance on legacy security approaches and tools will only hasten the increase in cloud-centric risks that already exist in most organizations.

Security executives must recognize the risks of such cloud initiatives are in many ways different from what they are accustomed to. They can no longer presume that their legacy vulnerability scanners, configuration managers, and identity management solutions can provide a holistic view into the highly volatile cloud environments being deployed today. They must fully embrace new solutions which have been designed and developed specifically for cloud platforms and find ways to integrate them with their legacy toolsets.

ANALYST REPORT

THE EVOLUTION OF EMAIL SECURITY PLATFORMS

DR. EDWARD AMOROSO

Three generations of email security approaches followed by the global community are explained starting with first generation focus on virus filtering, second generation focus on secure email gateways, and third (present) generation focus on cloud-based analytics. A simple evaluation framework is used to demonstrate how each subsequent generation introduced improved security coverage and accuracy. Predictions for future generation support of email security are included in the report.

1. BACKGROUND

The cyber security community understands fully the importance of email in offensive attack strategies by adversaries. The use of email vulnerabilities such as malware-laden attachments or phishing links to malicious sites remains an important component in attacks such as advanced persistent threats (APTs) initiated by offensive actors such as nation-states. For this reason, email security has emerged as a foundational component in the field of cyber security.

As a result, it is both educationally instructive and practically useful to understand the evolution of the email security threat from the perspective of both the offensive attacker and the defensive practitioner. This evolution spans the time starting with the first emails being sent in the 1970's over Arpanet to the present era, where email serves as the backbone for most business and even personal communications.[1]

The evolution is presented in three phases, starting with first generation use of computer viruses as payloads in emails sent from hackers to unsuspecting recipients. This is followed by a second-generation era in which gateways were used to filter and mitigate these viruses and other threats with varying levels of success. The present third generation of email security is shown to be characterized by advanced analytics that can intelligently secure email services.

2. FIRST GENERATION: VIRUS FILTERING

First generation email security emerged well-into the initial use of email as a personal and business platform. At the outset, it was unclear which threat models would drive users of email toward protection. Industry icon Bruce Schneier, for example, penned an early book on email security that focused on the use of public key infrastructure (PKI) for secrecy.[2] This encryption method never caught on and remains largely unused across email infrastructure.

Eventually, in the 90's and early 00's, the industry determined that viruses could be transported via email as a convenient mechanism for attack propagation. This required that two condition be met: First, the virus would have to execute in the targeted environment – and this was soon a non-problem as Microsoft Windows came to dominate the PC ecosystem for both business users and most home and family users.

But second, the virus attachment would need to be clicked on for download and execution. This requirement led to what we would now refer to as social engineering and phishing attacks. These were much simpler in the early days when users held the view that inbound email should be trusted. Recent advances in user awareness training have made phishing tougher, but still quite successful with many innocent and unsuspecting email recipients.

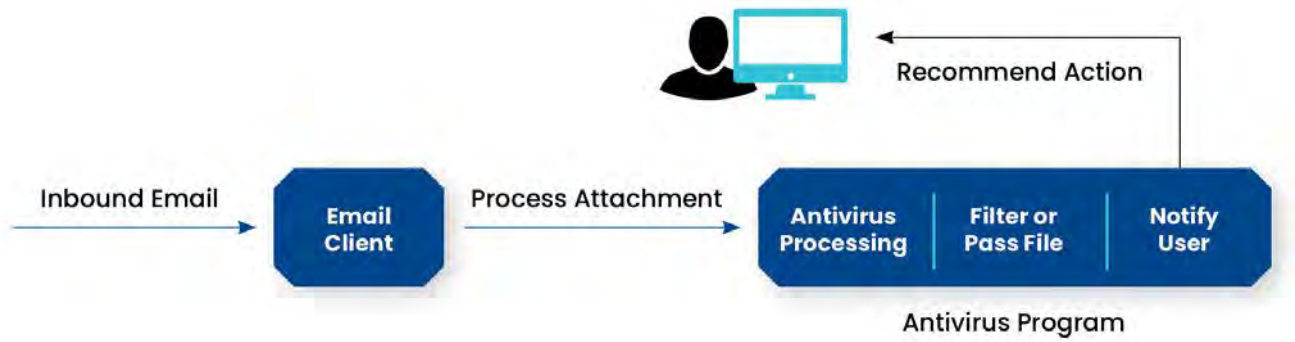

Figure 1. First Generation Email Virus Attacks

The most common solution implemented during this first generation of email security attacks involved using anti-virus software to secure the PC. The method relied heavily on the use of attack signatures to detect viruses. While this worked initially, variants were quickly developed to sidestep the detection.[3] Modern PC security tools use behavioral analytics and machine learning to detect the presence of unwanted malware more accurately.

Nevertheless, such methods presume the existence of viruses and focus on their detection and removal. This is a valuable strategy, especially when one assumes that exploits are mostly inevitable, and the cyber security industry has developed many solutions with this emphasis. Popular tools and platform that are postpend referenced with the detection-response (DR) designation involve "shifting right" to address an on-going threat.

3. SECOND GENERATION: GATEWAY PROCESSING

The second generation of email security involved attempts to be more proactive about these viruses and malware-leading URL links that would arrive in a recipient's inbox. The observation was made that if such threats come in through the usual series of store-and-forward nodes that characterize email transport, that one or more of these intermediate processing steps could be used for security inspection.

This is the origin of the well-known secure email gateway (SEG) platform, which has become almost ubiquitous with any enterprise business or government agency email system. The SEG was expected to be an extremely effective solution because control existed for how inbound email was handed. That is, the SEG could be placed in paths (via redirection) that would provide high levels of coverage for email being sent and received.

The good news during the 00's and early 10's was that this type of coverage, combined with processing methods that improved on early signature models, did offer (and continue to offer today) considerable risk reductions for inbound email attacks. This helps to explain why so many organizations continue to use a commercial SEG, and why this protection method is unlikely to disappear from enterprise architectures in the near term.

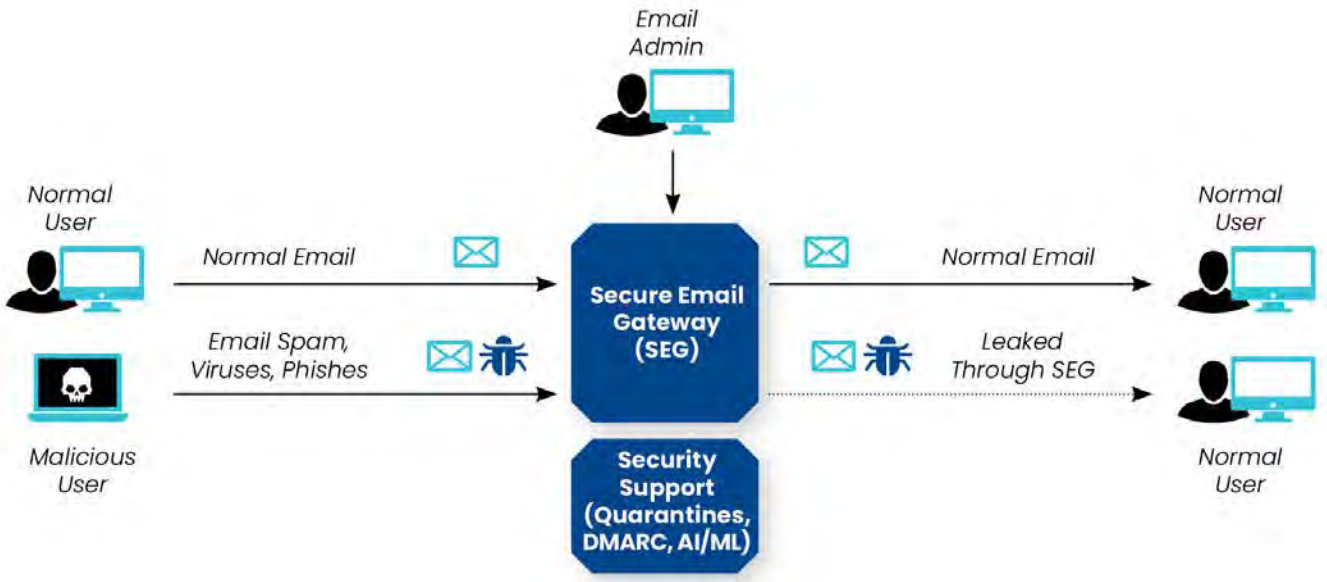

Figure 2. Second Generation SEG Filtering

Additional good news during this second generation was that standards teams developed new schemes for authenticating email sender identities. The open DMARC standard (Domain-based Message Authentication, Reporting, and Conformance) was the dominant contribution, and it allowed for senders of email to bind their originating IP address to any email carrying their domain.[4] This was designed to be done via DNS records which would allow recipients to enforce polices for handling mismarked email.

Sadly, two problems quickly emerged for DMARC during this second generation. First, many enterprise teams had trouble configuring records, especially in enforcement mode. The use of DNS TXT records for DMARC continues to make it tough for non-experts to navigate the difficult syntax and error-probe editing process required to properly set-up DMARC on DNS without causing unwanted negative side-effects.

Second, the DMARC standard has some awkward constraints that make it hard to use with cloud services. For example, there is a domain limit of ten rule sets in the Sender Policy Framework (SPF) portion of DMARC.[5] Since the whitelisting of public clouds will require several rule sets per service, many organizations will be forced to specify certain cloud sender security information by IP address – and this is neither convenient nor easy.

4. THIRD GENERATION: CLOUD-BASED ANALYTICS

The third and present generation of email security benefits from the pros and cons of anti-virus filtering, SEG processing, DMARC controls, and other capabilities used to reduce risk. This includes years of working with employees and users to help them make better decisions regarding security. This aspect

of conventional email security is particularly important because it highlights the synergy that can exist between systems and people at the human layer.

State-of-the-art platforms for email security today have precisely this attribute – namely that they can take full advantage of the things that software systems do best (e.g., process data) combined with things that humans do best (e.g., recognize patterns). The result is an analytics-based security approach natively tied to cloud infrastructure that has the strong potential to bend the risk curve downward for email infrastructure.

Elements of this third-generation email security solution include the following key protection features and risk controls:

- **Behavioral Analytics** – Behavioral analytics involve ingest of relevant factors, processing based on correlation and related strategies, and reporting in a variety of different means including via application programming interfaces (APIs) to other security tools.
- **Automated Learning** – With recent advances in machine learning algorithms, email security can improve continually based on patterns detected in test traffic or in live email traffic (e.g., for deep learning systems).
- **Personalized Protections** – Tailoring email security to match the preferences and usage patterns of individual users allows for more accurate handling and security. Some users might view an email as Spam, whereas another might view the same email as fine.
- **Cloud-Native Controls** – The use of cloud infrastructure has emerged as particularly useful for email security since it offers ubiquitous access for both ingest of threat intelligence, as well as for access to email systems.
- **Quantitative Risk Profiling** – Quantifying risk allows for effective reporting of email security posture, which can be helpful when stipulating minimum security levels or in measuring the benefits of a given security protection.

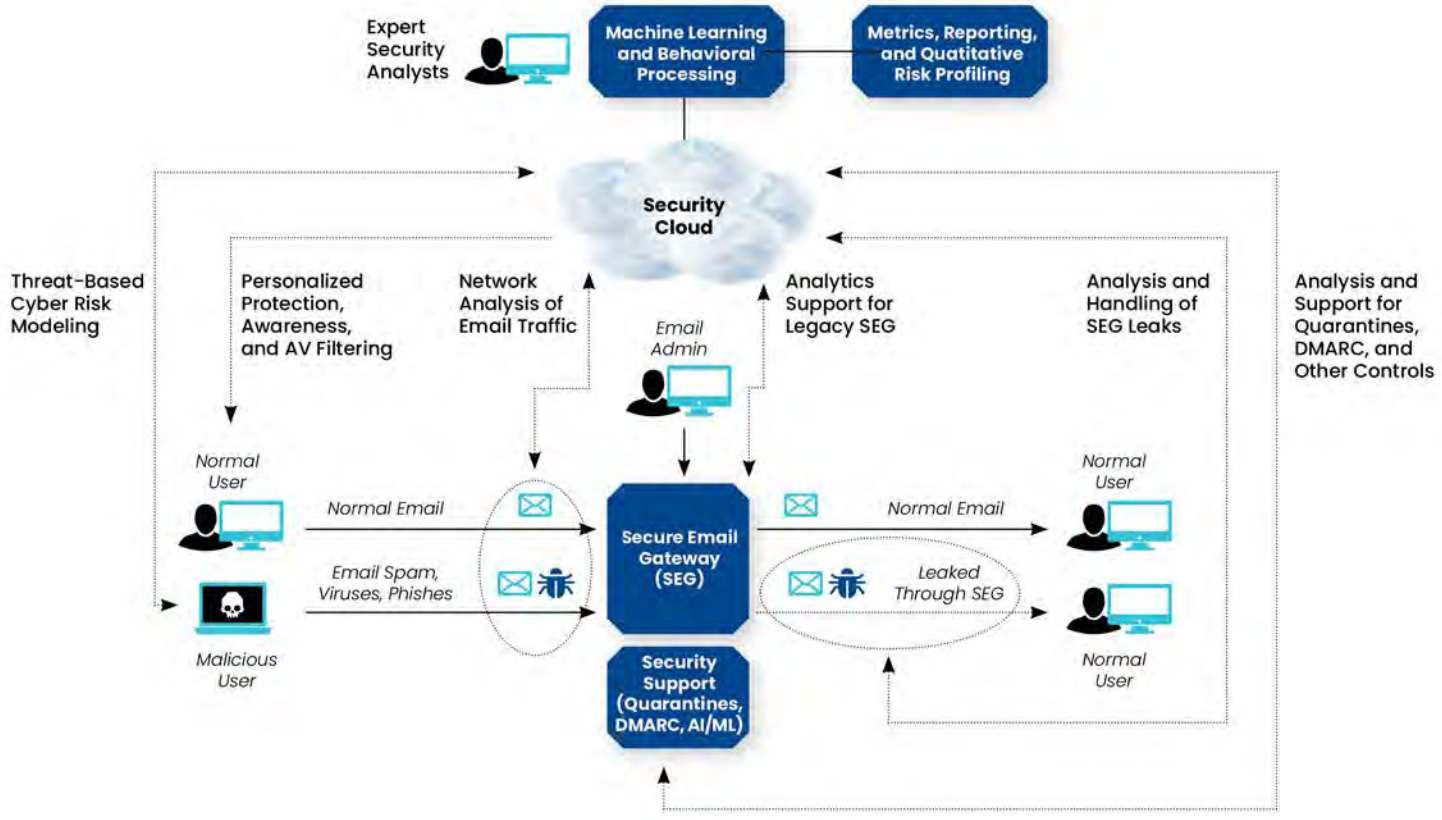

Figure 3. Third-Generation Analytics, Learning, and Advanced Controls

As each generation of email security progresses forward, the good news is that state-of-the-art solutions can incorporate the best practices and demonstrably useful elements of prior generation techniques. Nevertheless, even in the present generation of advanced analytic usage, business email compromise (BEC) and phishing attacks continue to occur, especially when combined with social engineering methods.[6]

The goal for email security will never be to reduce cyber risks to zero, but rather to address vulnerabilities sufficiently that email usage becomes a much lower concern for enterprise security teams as well as citizen users. It is an open question whether more intense attention to existing controls will be sufficient to achieve this objective, or if totally new security solutions will be required. The next section offers some views on this future state.

5. EVALUATION FRAMEWORK

To demonstrate how successive generations of email security have provided better handling and protection, it is helpful to introduce a simple evaluation framework. The objective is to identify the relevant aspects of email security that have changed over the years. These include the following attributes:

- **Email Threat Protection** – The purpose of email security obviously is to prevent, detect, or respond to threats – presumably with prevention as the ultimate objective. The good news is that successive generations of email protection introduce stepwise more effective security controls. The reason the problem remains however is that malicious activity has also increased and improved during the same period.
- **Transparency to Users** – An objective in any IT security control is transparency for users. This is especially true for email, since it is such a pervasive tool. During successive generations of email security, the obligations for users has increased, as evidenced by the extensive user awareness training typically required. Removal of such user friction should be an objective for future generation methods.
- **Lifecycle Costs to Organizations** – While license costs for email security tools have likely increased for most organizations, the corresponding costs for incident response have typically been reduced. This is true when suitable investments and good tooling have been put in place. The case to be avoided in the present generation involves high-cost email security tools without commensurate reduction in lifecycle costs.

	First Generation: Virus Filtering	Second Generation: Gateway Processing	Third Generation: Cloud-Based Analytics
Email Threat Protection	Weak Signature-Based Virus Detection	Improved Filtering and Behavior-Based Detection	Advanced Analytics-Based Detection
Transparency to Users	Users Filtered Viruses Manually	SEG Filtering Less Visible, but User Awareness Required	Low User Friction for Analytics, But Still High Awareness Need
Lifecycle Costs to Organizations	Virus Response Costs Began to Rise	Phishing Response Costs Reach Significant Level	Phishing Response Costs Started to be Managed

Figure 4. Effectiveness for Three Generations of Email Security

It is reasonable to conclude that great progress is being made in email security, as depicted informally in Figure 4 with the progression from red status (weak) to yellow status (improving) to green status

(effective). While phishing and business email compromise (BEC) are still problems, they usually stem from inadequate application of available tools and poorly conceived architectures – both of which will also improve in the coming years.

6. FUTURE GENERATION EMAIL SECURITY

The use of virus filtering, secure email gateway (SEG) processing, and cloud-based analytics has thus obviously reduced email security risk considerably. Such advances have been balanced, however, by malicious actors improving their own methods for targeting users. Automation has allowed for increased coverage in phishing attacks, even ones using individualized spear phishing methods. This has expanded the attack surface for email.

Based on the progression through three generations of email security and observations about trends in public cloud, SaaS, and networking, three observations can be made with respect to future generation email security. These observations should be viewed less as "predictions" and more as general "extrapolations" of on-going trends in how email security is likely to be handled in the next few years, post 2022.

- **Embedded Protections** – Email security protections should become more integrated natively into services provided by Microsoft, Google, and others. This does not imply that innovative technology companies working email security will cease to thrive and grow, but rather that the buying habits of end-users will demand that these new controls come pre-integrated and embedded into existing services.
- **Expended Intelligence** – Continued advances in artificial intelligence will lead to even more powerful controls for email security. Deep learning methods and computer vision, for example, will more than likely introduce new means for using live email streams as the basis for improving the accuracy and quality of the artificial intelligence models that serve as basis for the protections.
- **Increased Autonomy** – With the introduction of more autonomous computing methods (as evidenced in the trucking and automobile industries) will come greater confidence for user to rely on autonomous email assistants. Such assistants will reduce the monotony of handling routine email but will also come with advanced cybersecurity controls to avoid human errors.

The offense is also likely to improve its malicious techniques, and they should be expected to also rely on advances in artificial intelligence and autonomous computing to build more powerful attack tools. One would hope that the defense would be more aggressive in making progress and the shift toward intelligent autonomy should be particularly useful in reducing human errors, which have always been such an easily exploitable weakness.

REFERENCES

[1] Ray Tomlinson, "The First Network Email," openmap.bbn.com.
[2] Bruce Schneier, E-Mail Security, John Wiley & Sons, 1995.
[3] Timeline of Computer Viruses and Worms, Wikipedia.
[4] DMARC Website.
[5] DMARC Wiki.
[6] Jade Hill, "Inside the Business Email Compromise Problem," July 2021.

TAG CYBER

ANALYST REPORT

EXTENDED SECURITY POSTURE MANAGEMENT: AN OVERVIEW OF CYMULATE

DR. EDWARD AMOROSO

Modern enterprise teams demand continuous views of security posture to address ongoing cyberthreats and prevent enterprise drift, gaps and misconfigurations in today's dynamically changing enterprises. The result is a growing emphasis on extended security posture management—a protection approach exemplified by the commercial Cymulate platform.

INTRODUCTION

The requirement to understand security posture on an ongoing basis has emerged as one of the more challenging aspects of modern enterprise protection initiatives. Specifically, a discipline known as security control validation has emerged as a component of *extended security posture management (XSPM),* which offers detailed technical insights and prescriptive remediation assistance for practitioners, as well as high-level risk guidance for executives.

A key innovation found in most security-control validation systems involves the use of continuous testing and automation to generate accurate visibility and meaningful insights for action. Such automation enables the continuous optimization of controls and IT spending, while minimizing risk and helping assure the operational effectiveness of security systems. It also provides integration with existing cybersecurity components and infrastructure.

In this report, we introduce the concept of extended security posture management, with the goal of helping practitioners understand its relationship to related continuous protections. The commercial Cymulate platform is used to demonstrate the implementation of ongoing security-control validation in an enterprise context. Comprehensiveness, along with the ease of deployment and use, are its top design considerations, making this an attractive extended security posture management solution for working-level experts, senior-level managers and executive staff.

Furthermore, XSPM solutions allow CISOs to measure the effectiveness of security programs and maximize their return on investment, while also providing a clear understanding of enterprise risk levels and, more importantly, a prescriptive and prioritized list of how to further reduce risk. For the cybersecurity practitioner, XSPM solutions offer a clear-cut way to optimize enterprise cybersecurity, while understanding attacker tactics, techniques and procedures, as well as reducing enterprise attack surfaces and threat exposure.

EXTENDED SECURITY POSTURE MANAGEMENT

In the face of evolving threats, increased vulnerabilities and enterprise drift, XSPM solutions use attack simulations and other means to discover misconfigurations, gaps and vulnerabilities in order to establish attack feasibility and prioritize any risks that were found. When done on a continuous basis, enterprises can easily establish security baselines and trends over time, as well as receive automated and prescriptive technical remediation instructions. Executive reporting also includes industry benchmarking for companies that wish to compare their maturity and risk profiles.

A useful, high-level methodology for implementing XSPM includes three basic functional protection objectives: visibility into assets; optimization of posture; and the assurance of continued security. These three high-level management goals provide a helpful view of how XSPM platforms have been developed, and how enterprise teams can implement XSPM to offer the continuous improvement and validation of their deployed protection scheme.

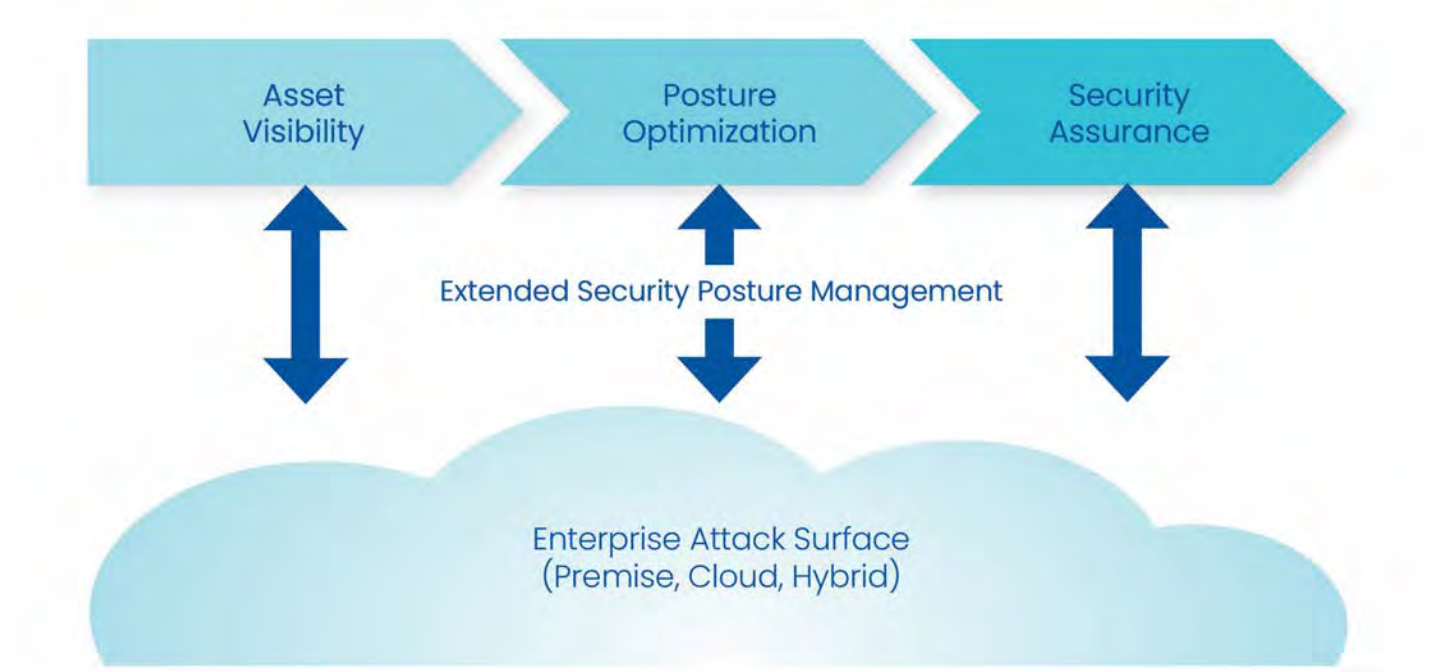

Figure 1. High-Level XSPM Objectives

The purpose of XSPM is to provide high confidence that an enterprise is being protected by properly deployed security controls that are configured without vulnerabilities and operating according to expected security parameters. The process is thus related to the emerging attack surface management (ASM) discipline and can easily complement or replace existing or planned security initiatives in this area.

XSPM is a valuable testing solution for Blue and Purple Teams seeking to test and optimize their first- and third-party security controls in order to find and remediate discovered security gaps and vulnerabilities. Advanced-scenario attack testing is a favored tool by Red Teams seeking to automate, scale and customize breach feasibility testing and threat-exposure management. Organizations of all sizes can take advantage of the Jump Start offering, which provides immediate security validation against new and top threats, while its Amplify solution provides a managed service offering an option for organizations that lack the manpower for security validation and ongoing assessment.

OVERVIEW OF CYMULATE

Cybersecurity start-up company Cymulate offers a commercial SaaS-based solution that effectively supports XSPM in accordance with the objectives outlined above. Cymulate security validation combines assessments of outside-in reconnaissance, security awareness, infrastructure resilience and security control validation in one platform.

Cymulate provides a range of security scoring, including baselines and trends over time, as well as the possibility to benchmark your scores against those of your peers. It also offers actionable remediation guidance for the following security management and support domains:

Security Control Validation
Cymulate initially provides validation of security controls via carefully designed, advanced attack simulations. This process results in the high confidence that controls for web, web application firewalls, email, applications, endpoints, segmentation and data loss prevention, as well as other resources, are all working as intended. The Cymulate platform receives frequent platform updates to ensure that the most recent adversarial indicators of compromises, attack techniques and procedures are integrated into simulated threat scenarios.

Using the MITRE ATT&CK framework and NIST 800-53, mapping, reporting and explanations ensure a common language that is well understood by the team. Prescriptive technical reporting ensures that any found misconfigurations and gaps are easily remediated. Cymulate has an extensive number of security control categories it can test, and it also incorporates third-party cybersecurity integrations, ensuring in-depth, accurate testing analysis and results.

Phishing Campaigns
Cymulate also allows you to run phishing campaigns to test your employees and provide clear-cut measurements of the risk found, which can lead to important, additional employee education and protection. Tied into Security Controls Validation, you can further see how your email infrastructure, operating systems, email gateways and endpoints handle the testing, as well.

Security Posture Reporting
Cymulate offers a comprehensive user interface for reporting detailed information about security posture. Reporting includes: drilling down into breach and attack simulation (BAS) results; continuous automated Red Team results; and advanced Purple Team data. Numeric scores help provide a rapid visualization of status, so that results can be interpreted not only by security practitioners, but also by management and executive teams (see Figure 2).

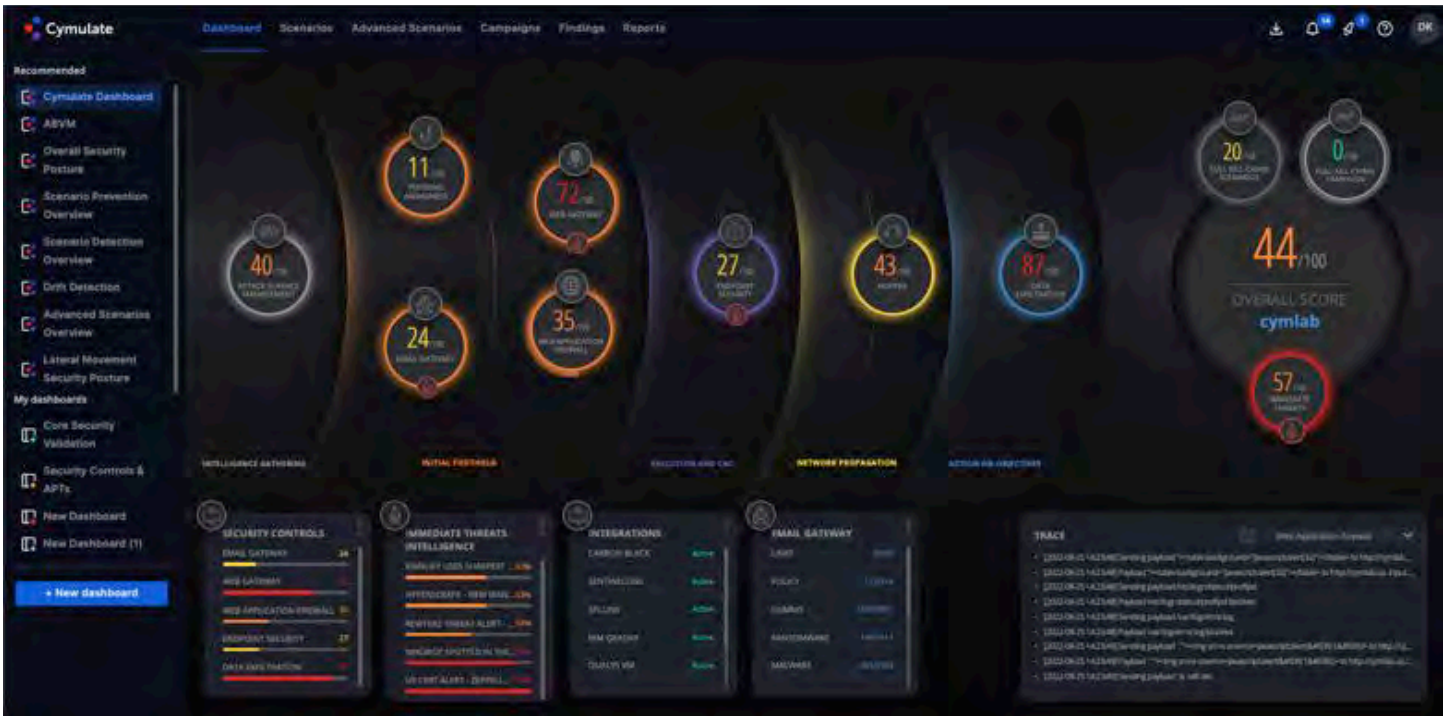

Figure 2. Cymulate XSPM Solution Dashboard
Showing Cybersecurity Posture Scoring Baselines and Trends

Attack Surface Management

The Cymulate platform identifies and prioritizes external-facing corporate assets, including web applications, exposed servers and other resources in order to expose the presence of exploitable vulnerabilities. Such attack surface management (ASM) is supported by the Cymulate Attack Surface Manager; it is valuable for complex environments that include a wide range of external assets, including shadow IT. An internal phishing awareness function provides complementary support.

Infrastructure Resilience

The Cymulate Lateral Movement component involves simulation of the East-West, lateral traversal, which is common in modern advanced persistent threat (APT) campaigns. The goal is to simulate the threat actions that result after an initial connection has been made to the target enterprise infrastructure. Misconfigurations and vulnerabilities are exposed via this traversal process, and support for risk measurement is included.

Continuous Testing

The Cymulate Continuous Automated Red Teaming solution involves support for attack campaigns against target enterprise infrastructure. The goal is to discover and expose vulnerabilities, weaknesses and soft spots in an enterprise, thus complementing the overall ASM solution. The testing process can be tailored to attack specific resources and can include tactics such as phishing to create initial access to target resources.

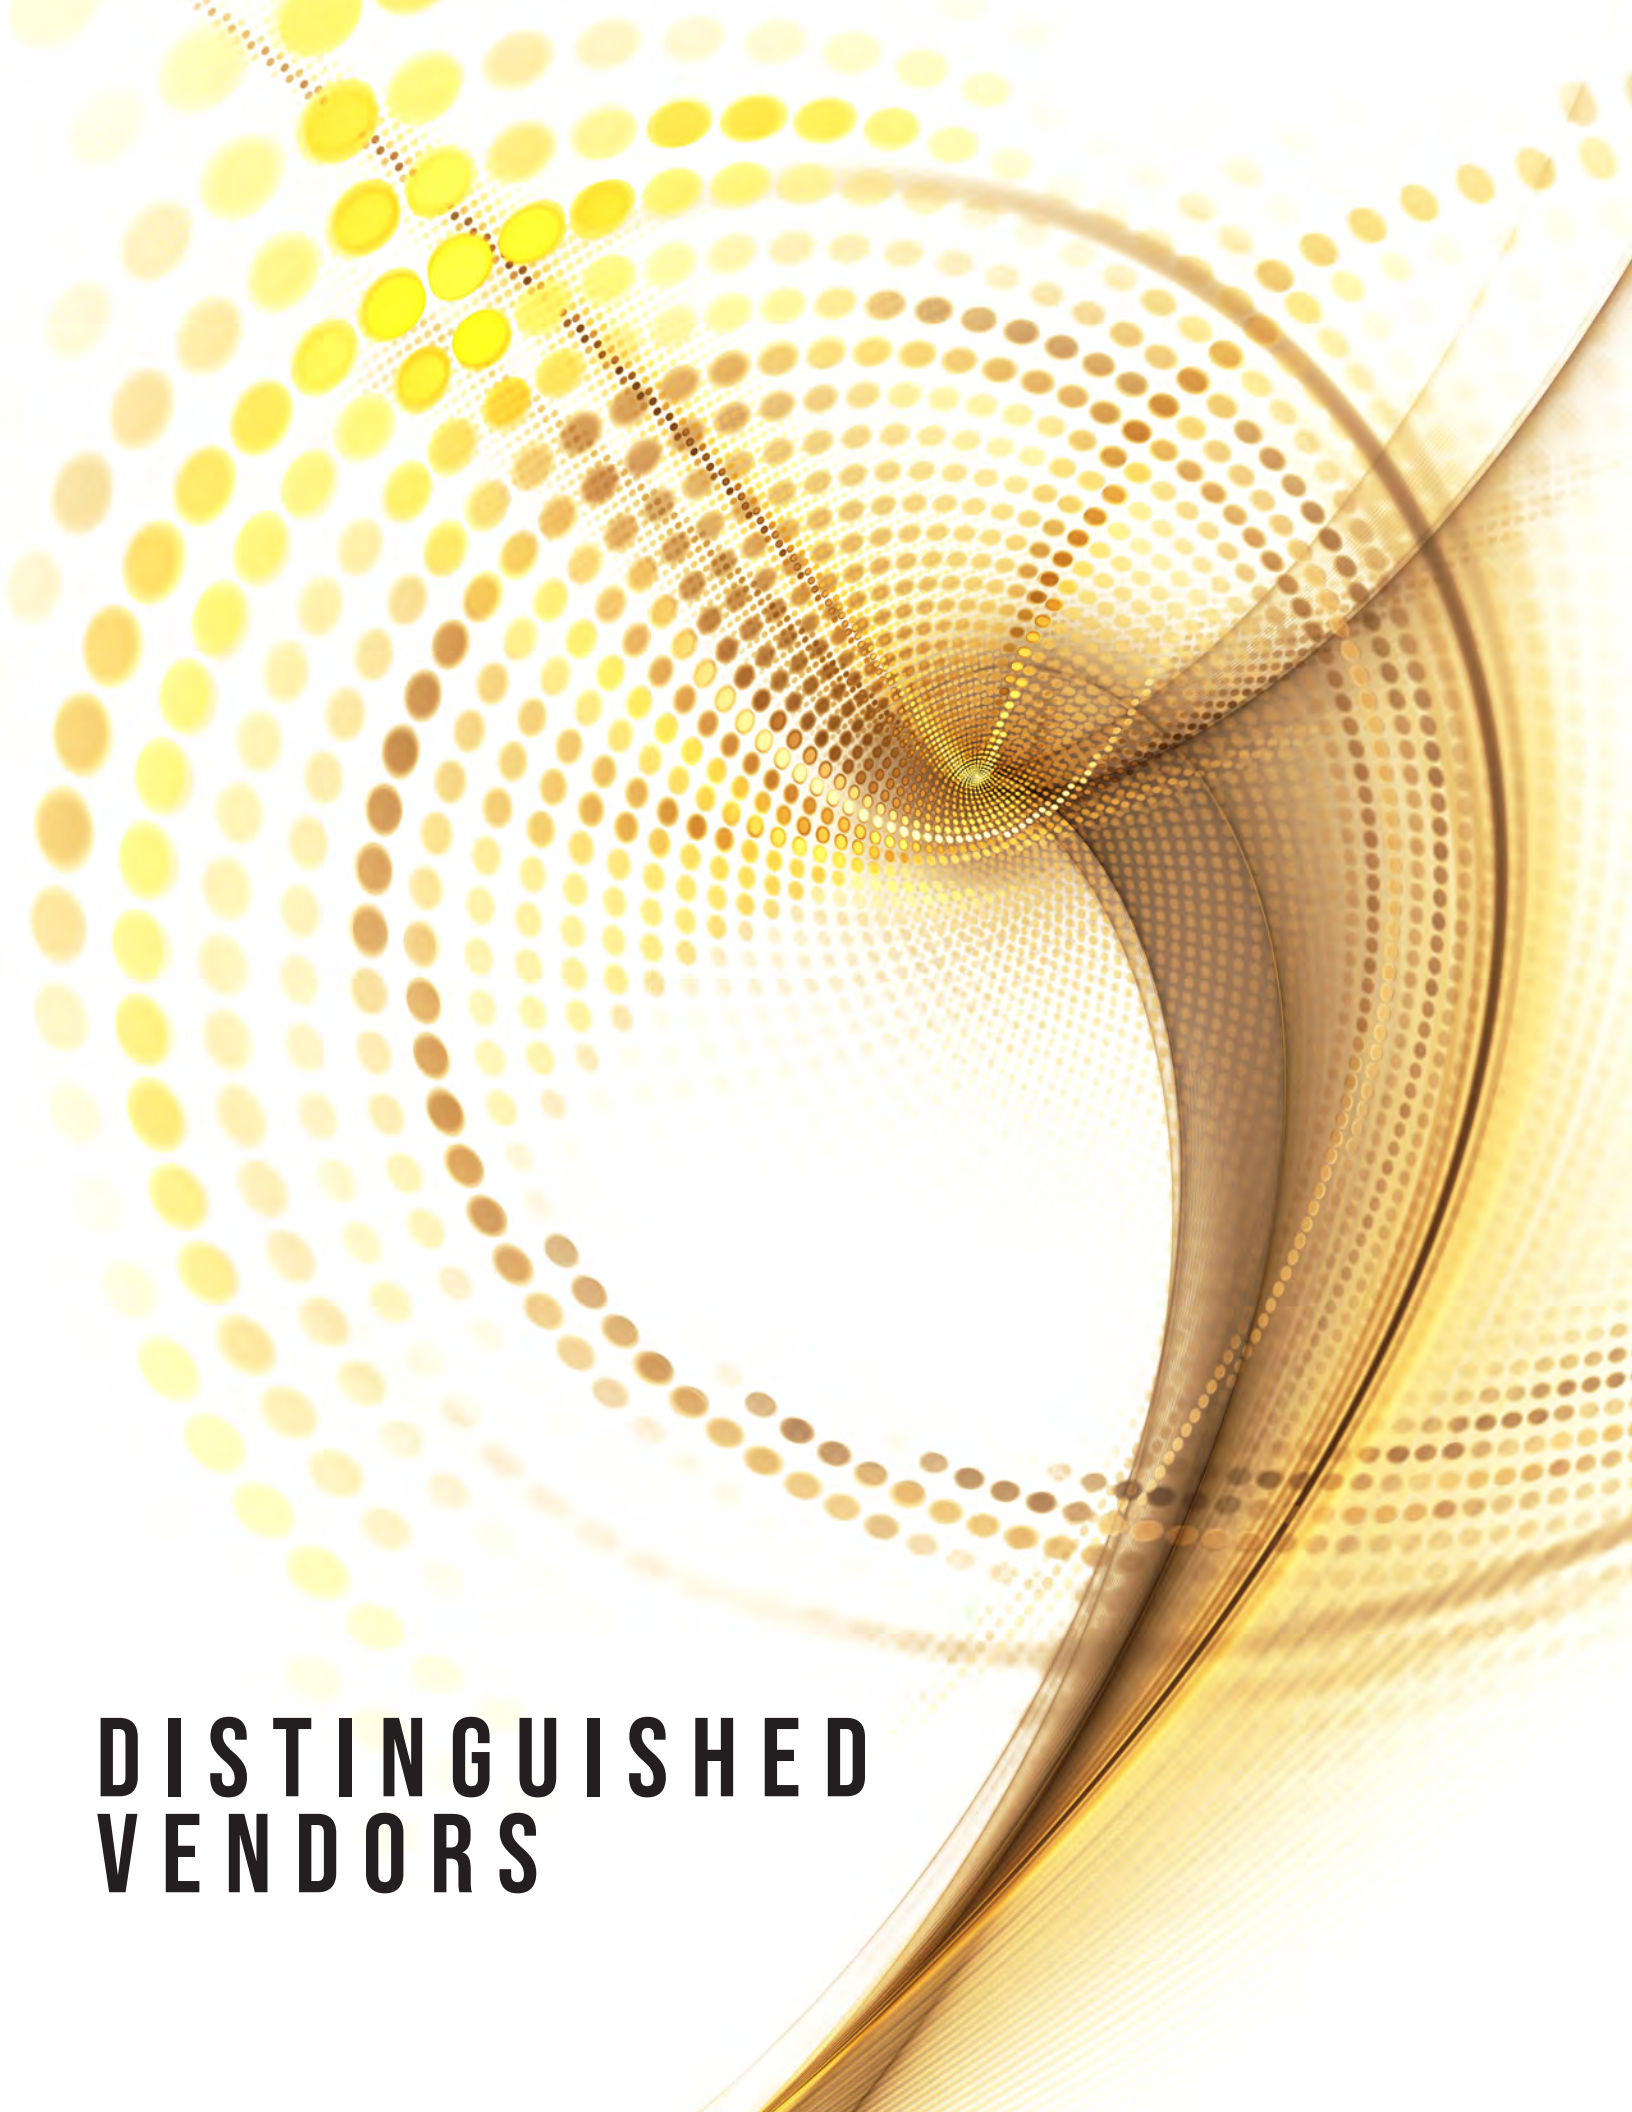

DISTINGUISHED VENDORS

TAG CYBER

DISTINGUISHED VENDORS

Q4 2022

Working with cybersecurity vendors is our passion. It's what we do every day. Following is a list of the Distinguished Vendors we've worked with this past three months. They are the cream of the crop in their area–and we can vouch for their expertise. While we never create quadrants or waves that rank and sort vendors (which is ridiculous), we are 100% eager to celebrate good technology and solutions when we find them. And the vendors below certainly have met that criteria.

allot

Allot is a global provider of leading innovative network intelligence and security solutions for service providers and enterprises worldwide. Its platform combines network-based security with home router and endpoint security to provide a unified security service for the mass market that's capable of protecting consumer IoT devices in the home, on mobile networks, and on public Wi-Fi.

ANOMALI

Anchored by big data management, The Anomali Platform, an Open XDR solution, drives detection, prioritization, and analysis to stop breaches and attackers in real-time. By fusing threat intelligence with precision detection capabilities, Anomali enhances threat visibility, automates detection, and accelerates threat investigation and response. The product suite includes ThreatStream®, Match™, and Lens™.

BLACKCLOAK™

BlackCloak extends enterprise security by protecting the personal digital lives of executives, Board Members, and high-access employees, and their families, from targeted cyberattacks and fraud. Their digital executive protection platform combines online privacy protection, personal device and home network security, and incident response, with a US-based SOC and concierge service.

CLOUD RANGE

Cloud Range cyber range training allows SOC analysts and incident responders to test and improve attack detection, response, and remediation capabilities within a safe environment. With virtual access or on-site training, users prepare for hyper-realistic attacks against their network and infrastructure and become better defenders.

TAG CYBER DISTINGUISHED VENDORS

2022

Concourse Labs offers a cloud configuration management platform with centralized, automated, Security-as-Code enforcement of security controls and policy. They enable enterprises to deploy mission critical applications to cloud with security, resiliency and regulatory compliance. Clients move away from point-in-time security snapshots and human-dependent security checklists to persistently secure, auditable processes and environments.

Cyber Security Solutions offers clients full protection and peace of mind with their all-in-one security solution, full compliance dashboard, secure file management system, 24/7 monitoring, industry certified practices and a personal onboarding process. Turnkey cyber solutions protect a variety of industries from insurance to law enforcement; medical to regulatory compliance.

Cymulate's Extended Security Posture Management allows organizations to measure and maximize operational efficiency while minimizing risk exposure. Based on real-time data, Cymulate protects IT environments, cloud initiatives and critical data against threat evolutions. Using simulation, evaluation and remediation, Cymulate empowers and defends organizations worldwide, including leading healthcare and financial services.

Cynamics guarantees unified network threat detection, providing a new cybersecurity paradigm. Combining AI and deep learning to analyze patterns and autonomously identify malicious behavior, Cynamics predicts threats long before they hit. Their patented AI technology—Novel Threat Detection—delivers 100% network coverage 24/7, reducing costs and complexities while removing onboarding roadblocks.

Elevate Security provides an open and extensible insider risk management solution designed to identify a company's riskiest users and prevent incidents before they adversely impact business. Elevate Security's platform integrates with leading technology systems and products to predict user risk and stop incidents before they start.

Fletch delivers instant answers to the most pressing cyber risk questions. Their Trending Threats app is like having a whole threat intel team in your back pocket, while their People Risk app enables you to investigate anyone in seconds.

TAG CYBER DISTINGUISHED VENDORS

2022

IGI The Cybersecurity People™

Headquartered in Rochester, NY with a remote workforce spanning across the United States, IGI delivers people-driven cybersecurity focused on individualized business strategy, enterprise-wide expertise, and unshakeable partnership. IGI is the OEM of the patented Nodeware® vulnerability management solution, an award-winning SaaS platform that continuously scans networks to identify critical vulnerabilities.

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.

Laminar

Laminar offers the first extensive cloud data security platform for everything built and run in AWS, Azure, GCP and Snowflake. The platform helps security and governance teams autonomously discover, prioritize and secure their data with continuous monitoring. Data is embedded into the cloud infrastructure, ensuring optimum defense against security breaches.

OptimEyes.ai

OptimEyes offers a unique AI-powered, SaaS-based solutions platform with fully automated risk frameworks to assist organizations by creating a single source of truth to manage their cyber, data-privacy and compliance risk. Risk models are customizable and provide an enterprise-wide view with real-time decision-making.

perimeter 81

Perimeter 81 is an enterprise-grade secure network platform that connects all users, in the office or remote, to all corporate resources: on premises, in public clouds, SaaS, or the open Internet. It is delivered as a cloud-native, simple-to-use service that is fully managed from a unified, single-pane-of-glass console.

PREVAILION

Prevailion is a cyber intelligence company that protects organizations by providing unmatched insights into real-time threats targeting their networks. Offering clients the Apex™ Platform that predicts pre-attacks, detects early stage infiltration and provides total supply chain visibility, Prevailion collects malicious communications originating from threats that have bypassed existing security controls.

TAG CYBER DISTINGUISHED VENDORS

2022

ReversingLabs unifies software development and security operations teams with its Titanium Hybrid-Cloud Platform for software supply chain security protection. The platform reduces attack surface risk by utilizing extensive intelligence monitoring to harvest thousands of file types at scale through deep software and file threat analysis, accelerating data release and response.

SailPoint is the leading provider of identity security for the modern enterprise, empowering organizations worldwide to put identity security at the core of their business. With a foundation of artificial intelligence and machine learning, SailPoint identity security delivers the right access to the right identities and resources at the right time.

Sevco Security offers persistent cybersecurity situational awareness for all corporate IT and Security Operations Teams. Comprised of cybersecurity leaders from top commercial vendors and U.S. Intelligence, Sevco Security is dedicated to giving enterprises all that's needed to ensure they know what everyone on-premise and off is doing at all times.

ShardSecure desensitizes sensitive data in multi-cloud, hybrid-cloud and private cloud environments while reducing management complexity and improving business continuity. Headquartered in NY with its engineering team in Sweden, ShardSecure offers innovative Microshard™ Technology that protects sensitive resources in the cloud by securing and preserving data backup and preventing file tampering.

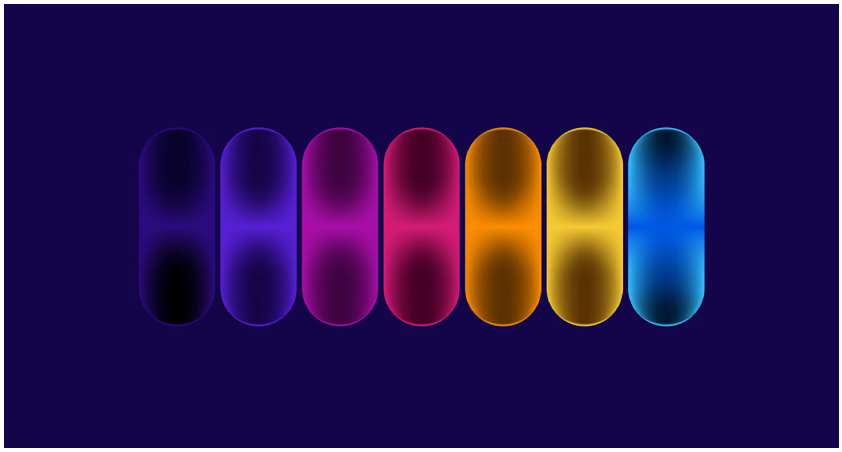

TAG CYBER DISTINGUISHED VENDORS

2022

Sicura is an automated security and compliance platform that seamlessly enforces and remediates technical security controls, fixes misconfigurations, and prevents security drift. The Sicura team of NSA veterans has built a product that bridges the gap between security and engineering teams, driving efficiency and improving enterprise security posture.

Sonrai Security offers total cloud security in one platform that unearths, prioritizes and removes risks across every part of the cloud. Their proprietary, big data analytics engine continuously updates the paths an identity has used or could use to access data, and offers visibility rooted in full context and actionable understanding.

Sphere is a woman-owned company that is redefining how organizations achieve controls across their environment. Its automation platform, SPHEREboard, provides an innovative approach that starts with collection and incorporates remediation of a client's most critical data, privileged accounts, and on-premises Messaging and Office 365 assets, while simplifying reporting and automating remediation to immediately reduce risk.

Sunday Security is a digital executive protection program, built to protect the world's executive teams beyond the enterprise perimeter. By harnessing our proprietary personal security platform coupled with our personal SOC, Sunday provides enterprise-grade personal cybersecurity to those at the enterprise who need it most.

Swimlane provides cloud-scale, low-code security automation for organizations of all industries and sizes. Our technology is rated as the #1 trusted low-code security automation platform. Our mission is to prevent breaches and enable continuous compliance via a low-code security automation platform that serves as the system of record for the entire security organization.

Sysdig is a software-as-a-service platform built on an open-source stack. Its Secure DevOps Platform provides security that lets clients confidently run containers, Kubernetes, and cloud services — allowing them to secure their build pipeline, detect and respond to runtime threats, continuously validate compliance, and monitor and troubleshoot cloud infrastructure and services.

TAG CYBER DISTINGUISHED VENDORS

2022

TITANIAM

Titaniam's advanced Data Security Platform utilizes encryption-in-use to make an enterprise's data immune to compromise without loss of functionality. The platform offers automatic compliance, flexible architecture, third party data control and fast, easy deployment. When all other security controls are breached, Titaniam continues to defend against ransomeware and other cyberattacks.

TRUEFORT

TrueFort offers application focused security at Wall Street speed and scale. Its award-winning fortress platform security system consolidates multiple fragmented security tools—cloud detection and response, zero trust segmentation, service account analytics, workload hardening and file integrity monitoring—to seamlessly shut down all unusual behavior and prevent malicious cyber infiltration.

vmware®

VMware is a leading provider of multi-cloud services for all apps, enabling digital innovation with enterprise control. As a trusted foundation to accelerate innovation, VMware software gives businesses the flexibility and choice they need to build the future.

VOTIRO

Votiro Cloud helps companies apply Zero Trust Content Security through its API-First Content Disarm and Reconstruction SaaS. With Votiro Cloud, enterprises can remove malware and ransomware threats in incoming files and content without using detection. Completely scalable and open to existing apps, data and security platforms, Votiro maintains instant content flows with no interruptions to productivity.

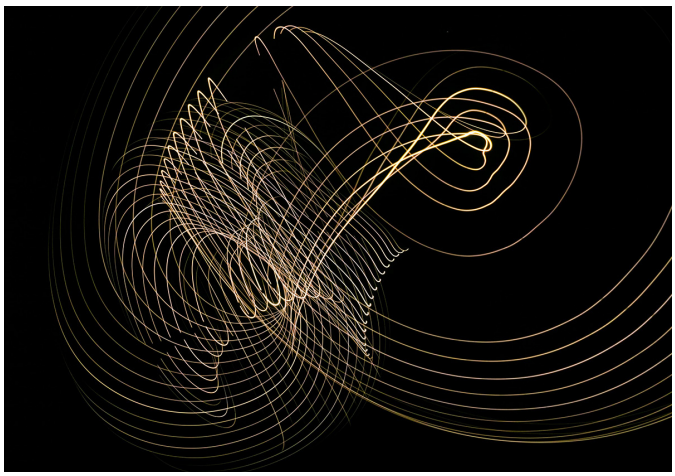

TAGCYBER