TAG Cyber Security Annual 2023

SPECIAL REPRINT EDITION

AUTOMATE YOUR CYBERSECURITY POSTURE WITH BALBIX

AN INTERVIEW WITH GAURAV BANGA, FOUNDER AND CEO, BALBIX

THERE'S MORE TO DEEPFAKES THAN MEETS THE EYE

SECURITY METRICS SOMETIMES MISS THE POINT

AUTOMATING CYBERSECURITY POSTURE ASSESSMENT: AN OVERVIEW OF THE BALBIX PLATFORM

TAG CYBER DISTINGUISHED VENDOR

Balbix

TAG CYBER 2023 DISTINGUISHED VENDOR

The need to reduce cyber risk has never been greater, and Balbix has demonstrated excellence in this regard. The TAG Cyber analysts have selected Balbix as a 2023 Distinguished Vendor, and such an award is based on merit. Enterprise teams using the Balbix platform will experience world-class risk reduction—and nothing is more important in enterprise security today.

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

REPRINTED FROM THE TAG CYBER SECURITY ANNUAL

TAGCYBER Balbix

AN INTERVIEW WITH GAURAV BANGA, FOUNDER AND CEO, BALBIX

AUTOMATE YOUR CYBERSECURITY POSTURE WITH BALBIX

Even with a wide variety of tools at their service, InfoSec teams that rely on manual workflows can no longer keep up with the ever-expanding enterprise attack surface. Networks can be compromised in an almost limitless number of ways, and these vulnerabilities open up businesses and organizations to serious damage.

Balbix automates cybersecurity posture by taking an accurate inventory of assets, while identifying the riskiest areas of the attack surface. It is geared to both mature and developing InfoSec programs in everything from start-ups to Fortune 500 companies. This scalable solution integrates with existing tools to reduce breach risk.

We met with Balbix to learn more about their AI-powered approach to cybersecurity.

TAG Cyber: Modern-day teams are drowning in cybersecurity data. How does your solution help them process this information overload to gain practical, useful insights?

BALBIX: Modern enterprises use dozens of cybersecurity tools, with each tool generating useful data about certain aspects of cybersecurity. Aggregating this data to produce a “big picture” of cyber risk has typically been done manually, often using proprietary algorithms and methods. Unfortunately, in recent years, this task has become untenable due to the exploding complexity of InfoSec programs. We must deal with different tool data formats and often inconsistent duplicates, as well as missing data about business context. The complex math required to calculate the next best steps for risk mitigation is nearly impossible. Furthermore, these aggregated models quickly become stale, because manual methods can’t keep up with constant changes in the threat landscape.

Our platform addresses this challenge by leveraging automation and AI. It continuously ingests and analyzes data from a company’s cybersecurity and IT tools to build a unified risk model. The system brings together data about vulnerabilities, threats, exposure, security controls and business criticality to prioritize security issues and surface the next best steps for risk reduction. The Balbix risk model is denominated in dollars (or other money units) and essentially maps from a digital/IT footprint to business risk. Security professionals can slice and dice their overall cyber risk in a variety of pivots—by business unit, attack vector, risk owner, etc.—and trace from dollars of business risk to the specific issues

Over a hundred machine-learning algorithms work together to normalize, deduplicate and correlate data to produce a unified picture of asset inventory and cyber risk.

driving risk. Our platform enables CISOs and their teams to make better cybersecurity decisions based on facts. An enterprise can build real-time cyber risk dashboards for business stakeholders, leading to the gamification of risk management. It also enables automated workflows for vulnerability management, which results in the faster mitigation of security risk issues. Ultimately, Balbix helps organizations drive increased efficiency, cyber risk reduction, cost avoidance and cost savings.

TAG Cyber: How does Balbix assist in automating vulnerability management?

BALBIX: With our solution, organizations can maximally automate workflows for identifying and prioritizing vulnerabilities, by dispatching these issues to risk owners and then driving mitigation and verification. To automate vulnerability assessment, Balbix maintains a comprehensive, real-time asset inventory and software bill of materials for the enterprise. This information is continuously evaluated against vulnerability data provided by software vendors, government sources and researchers to identify and tag vulnerable assets. Our platform automatically maps vulnerabilities to TTPs and continuously tracks real-world threat information. For each vulnerability instance on every asset, Balbix evaluates the effectiveness of security controls against these TTPs, as well as the business criticality of the asset to determine priority. Our platform also provides specific patch/fix information and other context to support mitigation efforts by relevant risk owners. If stakeholders choose to accept risk for some issues, then Balbix tracks this information. With Balbix, organizations can calculate and configure appropriate service level agreements (SLAs) for vulnerability management, based on their risk appetite and tolerance. Companies can build dashboards and reports to track/trend SLA compliance and cyber risk for each risk owner, asset type, application and business unit—geo, as well as the overall enterprise.

TAG Cyber: Tell us about the benefits of your Asset Inventory dashboard.

BALBIX: Our asset inventory dashboard provides organizations with a comprehensive and real-time view of the enterprise's asset inventory and software bill of materials. The Balbix data model includes over 450 distinct asset attribute types, all of which are surfaced in our asset inventory views. In addition, applications are mapped to the corresponding infrastructure asset, and each asset is tagged with relevant business context. With our asset inventory, security and IT professionals have the accurate, comprehensive information that is needed in their daily tasks. They save time that otherwise would be needed to follow and correlate information across multiple tools. There is no need to export and analyze data in Excel while solving problems, validating compliance or reporting. Overall, this saves up to hundreds, sometimes thousands, of hours of effort. Perhaps most importantly, Balbix Asset Inventory provides more than just visibility; it is tightly integrated into other Balbix capabilities that deliver maximally automated risk prioritization and mitigation.

TAG Cyber: What is the "Balbix Brain" and how does it help companies use AI to stay ahead of cyberattacks?

BALBIX: The Balbix Brain continuously ingests data from enterprise cybersecurity and IT tools, as well as external data sources. Over a hundred machine-learning algorithms work together to normalize, deduplicate and correlate data to produce a unified picture of asset inventory and cyber risk. The system brings together data about vulnerabilities, threats, exposure, security controls and business criticality, as well as performing probabilistic math calculations for cyber risk—asset by asset, application by application, and group by group across the enterprise. Unlike other AI platforms, the Balbix Brain was specifically designed for the model to explain itself and support traceability from dollars of risk to drivers of risk. As the complexity of the enterprise attack surface increases, cybersecurity data analysis becomes increasingly difficult. Balbix Brain provides critical capabilities for organizations to understand their gaps and associated risks, and close these gaps before adversaries can cause damage.

TAG Cyber: What are the top cyber risks facing companies in 2023?

BALBIX: There are three drivers making cybersecurity in 2023 more challenging than before. First, there is *AI/ML powered innovation in cyberattacks*. For example, the AI chatbot, ChatGPT, is already being used to generate very sophisticated phishing attacks. Most organizations are completely unprepared for automated AI-powered cyberattacks. Next is *flat or reduced cybersecurity spending*. Many InfoSec teams are facing budget cuts for tools and people due to the poor macroeconomic outlook. Unless organizations make a concerted effort to do more with less by leveraging more automation, they will face sharply higher cyber risk. Finally, there are the factors of *hopelessness, indifference and hubris*. Will your current InfoSec setup—people, processes and tools—deliver in 2023? Now is the time to take a step back and review if you have a good handle of your attack surface, and how your mean time to mitigate (MTTM) risk stacks up against the adversary's key metric—i.e., less than 15 days to weaponize newly found security vulnerabilities. Do all your stakeholders understand the amount of cyber risk you have on the books in dollar terms, and are they engaged actively in risk management? You may not like the results of your review, but now's the time to act!

THERE'S MORE TO DEEPFAKES THAN MEETS THE EYE

DAVID HECHLER

TAG CYBER/DALL-E

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

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

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

THEY'RE NOT ALL BAD

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

Kim Jong-un deepfake

REPRESENTUS

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

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

Joaquin Oliver deepfake

CHANGE THE REF

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

Then there's Bruce Willis, whose health problems led him to retire from acting. But he recently **made a deal** to allow a company called Deepcake (that's not a typo) to map his face onto the body of another actor for a **commercial**. Though there was some disagreement about the circumstances, the message Deepcake was announcing was clear. As was the company's aim to launch a new industry. Actors who can no longer act, the company seemed to be saying, or actors who have a commitment to perform that conflicts with another opportunity elsewhere, can now digitally clone themselves by authorizing deepfakes.

GRAY AREAS

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

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

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

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

THE DARK SIDE

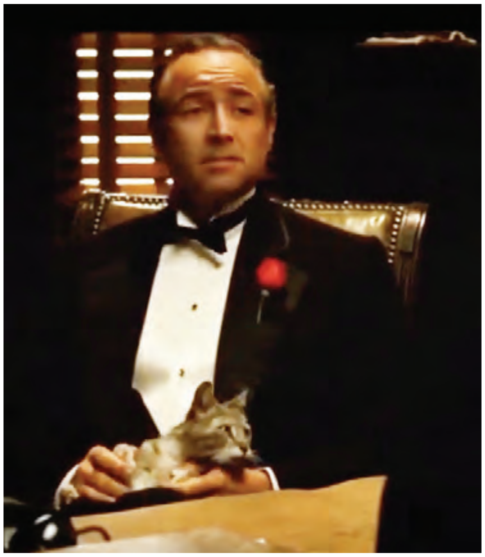

Nicolas Cage as Marlon Brando deepfake

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

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

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

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

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

PLAYING FOR HIGHER STAKES

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

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

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

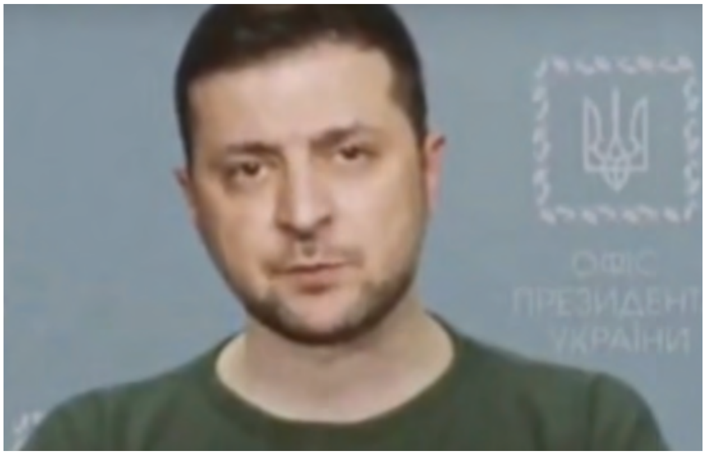

Volodymyr Zelensky deepfake

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

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

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

HELP NOT WANTED

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

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

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

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

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

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

SECURITY METRICS SOMETIMES MISS THE POINT

JOHN J. MASSERINI

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

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

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

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

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

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

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

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

OPERATIONAL VS. RISK METRICS

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

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

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

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

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

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

Ponder that for a moment. Most of us deal with boards and management teams that comprise scores of participants. Your metrics need to make sense not to the one person you are speaking to, but the dozen or more board members who come from diverse backgrounds and experiences. You don't have one different map of the world to deal with, but dozens—dozens of people who all heard the exact same words you spoke, and who all interpreted those words slightly differently. Well-planned metrics bridge the communications gap that comes with having multiple world maps in your boardrooms.

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

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

OPERATIONAL METRICS

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

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

RISK METRICS

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

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

- **Vulnerability Metrics:** The stalwart of the metrics world is undoubtedly reporting vulnerabilities. The key to effective vulnerability metric reporting is to relate them to the potential financial impact on the company. Do not report to the board a count of generic five-tier risks (none through critical) without offering insight into the financial impact of your critical systems. Again, avoid using these numbers to instill fear, but rather, put these findings into context by associating them with the revenue that could be impacted by attacks.
- **Identity Metrics:** As more enterprises begin planning their long-term, zero trust initiatives, having a clear understanding of your access controls is critical. Understanding how identities and accounts are created, maintained and ultimately deleted is a foundational necessity when you consider zero trust. Tracking topics such as role ratio, mean time to close, recertification requirements and "out of compliance" metrics will drive a deeper understanding of identity-related risk throughout the enterprise. Also, do not forget to collect and evaluate identity metrics around your AWS/GCP/MSA cloud environments, as access control risks are substantially more risky when you consider most DevOps processes.
- **Availability Metrics:** It seems all too often the availability of a system is prioritized well behind the confidentiality or integrity of a system, rather than giving it an equal footing. Have you done a business impact analysis on that 30-year-old system that runs that old Cobol-68 program which just happens to drive 75% of your revenue? Well guess what? The board wants to know you're on it and there's a plan to ensure it's upgraded, migrated or backed up even though there isn't a published exploit anywhere in the world. If you've forgotten what C.I.A. (confidentiality, integrity and availability) is perhaps it's time for a refresher.
- **Regulatory Metrics:** We all have them—whether it's PCI, HIPAA, SOX or any other government/industry related acronym—and regulatory requirements are something we all have to deal with. When discussing these risks with your board, do not just talk about the gaps you have. Make sure you also articulate the potential fines—especially in this GDPR world—and how those gaps could directly impact the levels of fines faced. Again, it's easy to fall into the trap of instilling fear with this, but try to avoid it. Use as much realistic data as possible, especially when dealing with publicly disclosed fines.

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

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

AUTOMATING CYBERSECURITY POSTURE ASSESSMENT: AN OVERVIEW OF THE BALBIX PLATFORM

DR. EDWARD AMOROSO

Establishing cybersecurity posture is an important step toward mitigating the cyber risks to an enterprise. Automation is the best approach for such assessment—one that builds on existing foundational security methods. The Balbix[1] Security Cloud is shown to automate this cybersecurity posture assessment process effectively.

INTRODUCTION

A major goal for enterprise security teams is to identify the attack surface that malicious adversaries can exploit. Such identification is the first step in mitigating cyber risk, and while the process might be simple to define, it is much tougher to implement. Modern enterprise infrastructure typically includes a complex mix of on-premises, cloud, SaaS, and hybrid infrastructure connected via proprietary and off-the-shelf software apps.

The process of defining all relevant vulnerabilities (or lack thereof) for a given attack surface is often referred to as the *security posture*. As one might expect, this has traditionally been achieved using a combination of scanning tools, asset databases, penetration test results, and other security tool output. Aggregation of this data has typically been done manually, often using proprietary algorithms and methods.

In this report, we explain how cybersecurity posture assessments can be automated. This is an important objective because it can establish a more continuous view of posture and will greatly reduce the possibility for coverage or completeness deficiencies. The commercial *Balbix* platform is used to illustrate how such a practical, automated assessment can be done in an enterprise context.

SECURITY POSTURE FOUNDATIONS

The challenge of establishing security posture can be visualized by mapping the assets of an organization against potential attacks. The two-dimensional structure that emerges is further complicated by the consequences, expressed in terms of financial loss,[2] that can result from a compromise. The result is a three-dimensional structure with a massive number of asset-attack-consequence mappings.

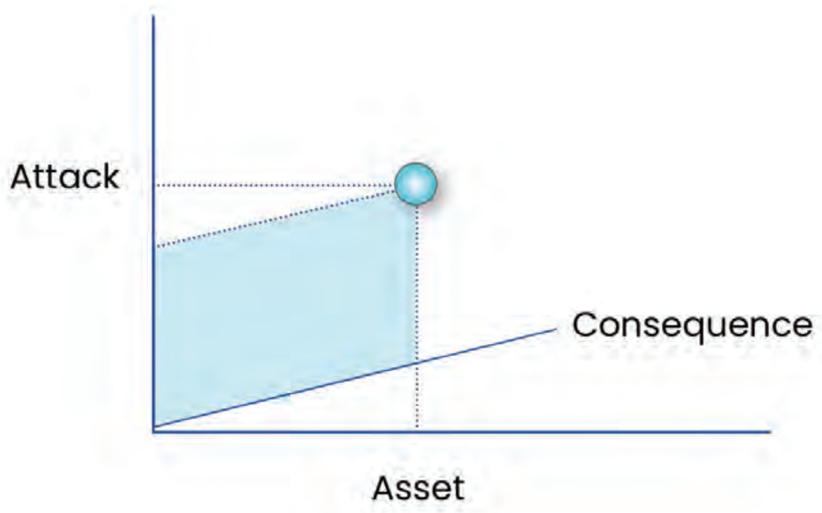

Figure 1. Mapping Assets, Attacks, and Consequences

The goal of gaining visibility into the present and ongoing status of cybersecurity controls is obviously not new. The primary means by which this goal has been addressed in the past includes familiar methods, many of which remain useful, but none of which have properly met the challenge. Since these traditional methods play a role in more evolved strategies for posture assessment, it is worth briefly reviewing the benefits of each.

Breach Simulation
One way to demonstrate the effectiveness of internal controls is to test them continually. To that end, so-called breach and attack simulation (BAS) tools have emerged to help enterprise teams determine the effectiveness of deployed security systems and tools. BAS implementations typically involve placement of active agents on either side of a control to continually test its ability to block attacks.

The advantages of a BAS approach include automated operation and continuous coverage. The disadvantages include limited flexibility and difficulty expanding to include more complex attack campaigns. Ultimately, BAS solutions are likely to find their way into a target security architecture, either as stand-alone platforms or as functional components of a more comprehensive protection architecture.

Vulnerability Scans
An additional major aspect of security posture assessment involves scanning networks, systems, and other resources for evidence of exposure. Operating a security scanner is perhaps the most familiar and traditional aspect of vulnerability detections and, as such, it is not only a requirement in every framework, but is also a major expectation of executives, board members, and other influencers.

The primary advantage of vulnerability scans is the familiar, mature data output that can support existing security and compliance programs. Most participants in enterprise security expect and understand this data, so scanning is essential in this context. The primary disadvantage is that scan data is prone to gaps in coverage and significant misinterpretation by executives and other stakeholders.

Penetration Tests
Penetration testing is also an effective means for identifying security vulnerabilities, especially ones that are subtle and not easy to find. For many years, enterprise security teams have relied on expert white hat hackers to probe, scan, and explore visible infrastructure with the goal of finding exploitable errors before a malicious adversary might find them and cause real consequences.

The advantage of penetration testing is that it is good at identifying the presence of security issues. That is, in environments where it is not generally accepted that exploitable holes exist, penetration testing can provide clarity. The biggest problem with penetration testing, however, is that it is an insufficient means for demonstrating the absence of problems. Not finding something during a penetration test doesn't mean that it doesn't exist.

Crowdsourced Testing
Finally, the use of vetted hackers (e.g., bug bounty) to help identify vulnerabilities has been an important component of an enterprise security posture assessment program. Since techniques, skills, and insights can vary so much between expert testers, having a large group of such individuals targeting a given system is a major advantage that offers depth of coverage and scope that cannot be reached by an individual.

The advantage of crowdsourced testing is the wide range of skills that can be harnessed to identify exploitable vulnerabilities. A drawback, however, is that considerable time and effort is required to properly vet and manage the ethical hackers. This workload can be mitigated through partnership with a capable commercial vendor, but it nevertheless represents a considerable hurdle.

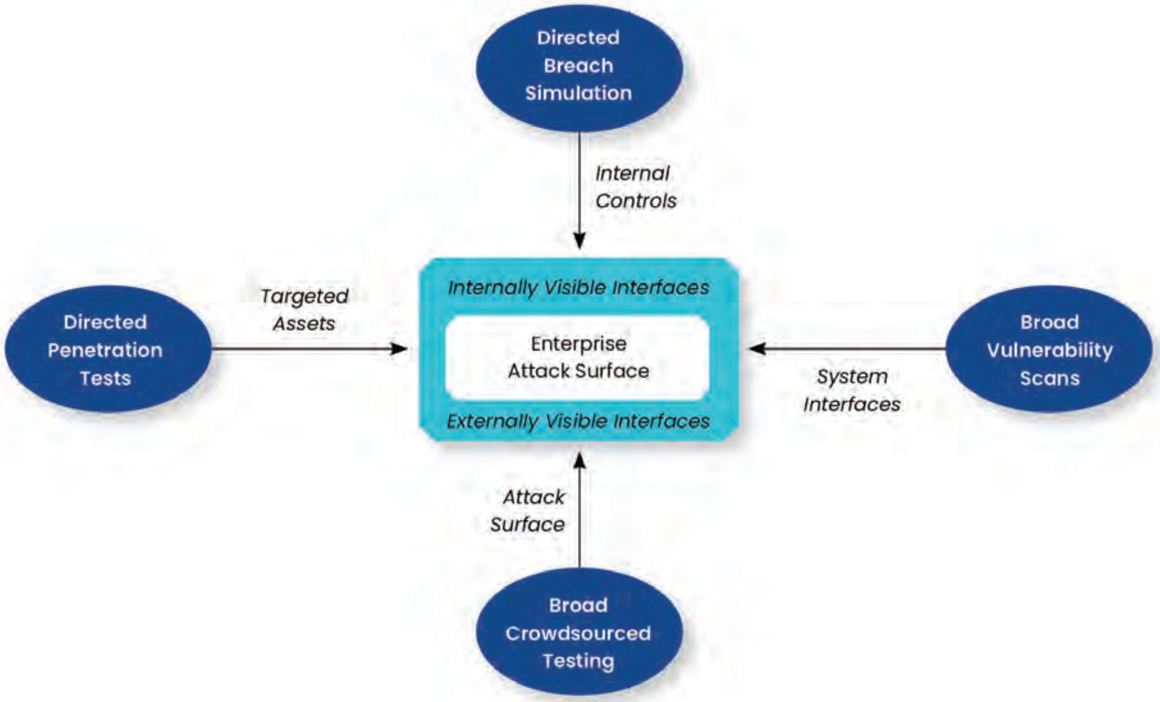
Directed Breach Simulation
Internal Controls
Targeted Assets
Directed Penetration Tests
Internally Visible Interfaces
Enterprise Attack Surface
Externally Visible Interfaces
System Interfaces
Broad Vulnerability Scans
Attack Surface
Broad Crowdsourced Testing

Figure 2. Common Traditional Methods for Identifying Security Posture

The challenge with these various methods is that while they each provide some degree of visibility into security posture, they remain disparate and uneven in terms of their automated or manual control. In the next section, we introduce a commercial platform from Balbix that uses automation as the basis for establishing an accurate, scalable view into the security posture of an organization.

CASE STUDY: BALBIX APPROACH TO AUTOMATED SECURITY POSTURE

The commercial Balbix platform provides for cybersecurity posture automation. It was created to complement existing vulnerability management and related security posture capabilities deployed into the enterprise, while also addressing the major challenges and shortcomings that such functions have typically exhibited in practice for most security teams. Some teams will find that Balbix can replace their existing posture tools.

Automated Asset Discovery
The first goal of the Balbix platform is to address the ongoing challenge of inaccurate and incomplete asset inventories. Without clarity around the specific devices, apps, endpoints, and other resources in use across the enterprise, it becomes impossible to have a complete measure of security posture. This challenge is further driven by the consistent change that occurs even for those assets for which an inventory has been established.

Balbix addresses this requirement through automated, continuous monitoring of the enterprise, including traffic flows, to discover assets. The types of assets that emerge from this task include on-premises and cloud-based devices, applications, systems, and services, including managed and unmanaged assets. Fixed and mobile systems, including Internet of Things (IoT) devices are also included in the asset discovery capability.

Data is discovered in the Balbix platform using a library of connectors that can handle two primary scenarios: *streaming connector*-based collection of data in motion, and *snapshot* connector-based collection of data at rest. Both take advantage of available interfaces including data dumps and application programming interfaces (APIs) to ingest the data necessary to build accurate inventory views.

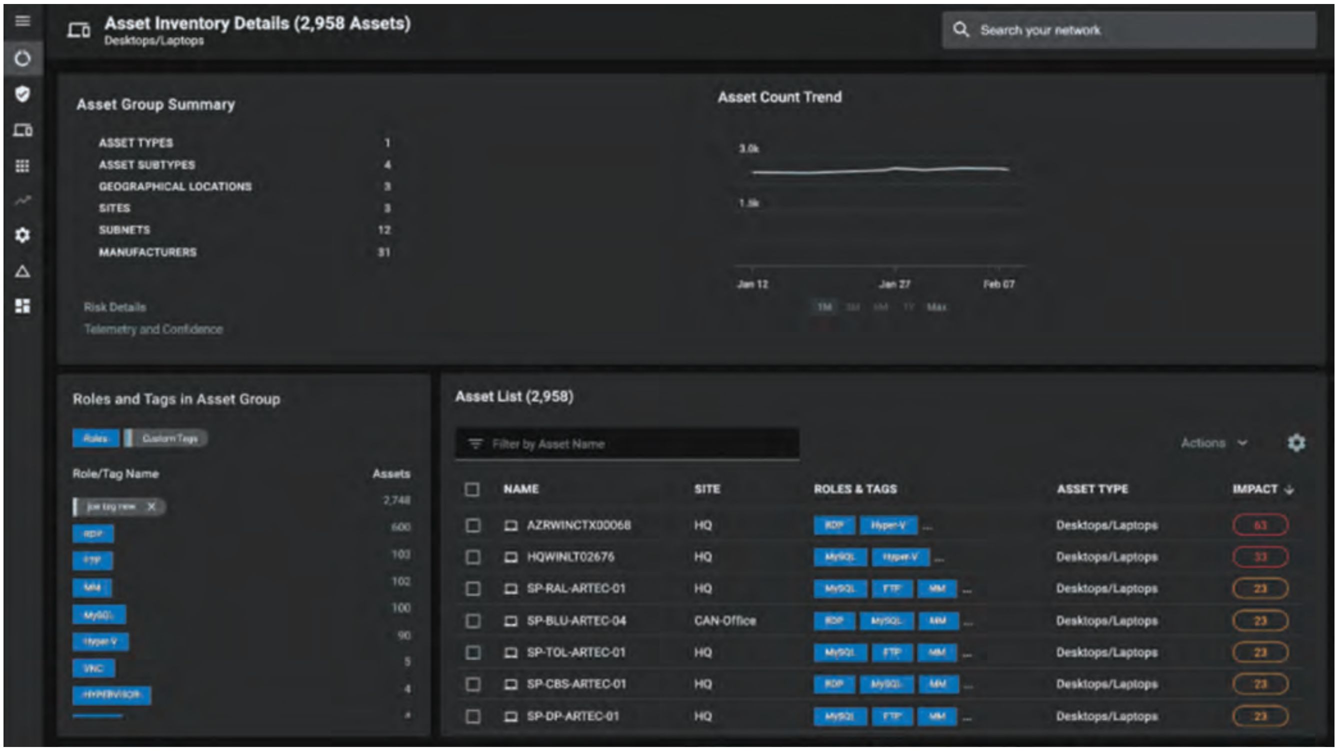

Asset Inventory Details (2,958 Assets)
Desktops/Laptops

Asset Group Summary

ASSET TYPES	1
ASSET SUBTYPES	4
GEOGRAPHICAL LOCATIONS	3
SITES	3
SUBNETS	12
MANUFACTURERS	31

Risk Details
Telemetry and Confidence

Roles and Tags in Asset Group

Role/Tag Name	Assets
joe tag new	2,748
RDP	600
FTP	103
MM	102
MySQL	100
Hyper-V	90
VNC	5
HYPERVISOR	4
[illegible]	4

Asset List (2,958)

Filter by Asset Name

□	NAME	SITE	ROLES & TAGS	ASSET TYPE	IMPACT ↓
□	AZRWINCTX00068	HQ	RDP Hyper-V ...	Desktops/Laptops	63
□	HQWINLT02676	HQ	MySQL Hyper-V ...	Desktops/Laptops	33
□	SP-RAL-ARTEC-01	HQ	MySQL FTP MM ...	Desktops/Laptops	23
□	SP-BLU-ARTEC-04	CAN-Office	RDP MySQL MM ...	Desktops/Laptops	23
□	SP-TOL-ARTEC-01	HQ	MySQL FTP MM ...	Desktops/Laptops	23
□	SP-CBS-ARTEC-01	HQ	RDP MySQL MM ...	Desktops/Laptops	23
□	SP-DP-ARTEC-01	HQ	MySQL FTP MM ...	Desktops/Laptops	23

Figure 3. Balbix Platform – Discovered Asset Details

Continuous Cybersecurity Asset Management

Once a complete picture of security posture has been created for the entire attack surface, the obligation emerges to manage and maintain the asset inventory and associated context in a unified manner based on automated platform support. The Balbix platform includes support for vulnerability and risk management workflows to ensure that assets are managed continuously to provide accurate security posture even as the attack surface evolves.

The collected data is used to categorize and manage assets based on their visible attributes, including internet protocol (IP) addresses, domain name system (DNS) information, and other signals that can be used to identify entities. The technique used by Balbix to normalize the accurate asset inventory view is called *host enumeration logic*, which supports stateful, intelligent de-duplication, sanitization, and other data clean-up tasks.

Such tasks must be performed at all levels of the technology stack, each of which will provide a different type of asset-related information. Layer 7 analysis, for example, will be useful to extract application-level information about assets, whereas layer 3 and 4 analysis will be useful to extract information about packet headers and protocol behaviors. The goal is to combine this collection into a unified view of the discovered asset.

Risk-Based Vulnerability Management

A major problem reported by enterprise teams is the large volume of alerts that is collected by typical vulnerability management and scanning tools. It is common for the number of alerts to become so high that security teams cannot maintain proper categorization, handling, and mitigation. This situation is ironic, because the success of vulnerability management programs is often measured based on the numbers of alerts generated.

The Balbix platform handles the volume of vulnerability management by ingesting and analyzing data from a massive number of security-related sources. These sources include vulnerability assessment tools, security scanning platforms, threat and vulnerability feeds, BAS tools, penetration testing results, crowdsourced security test output, endpoint controls, and more.

Enterprise Vulnerability Prioritization

Prioritizing vulnerabilities requires attention to relevant factors, most of which will vary in intensity between environments. The Balbix approach involves establishing five major categories of factors—severity, threat, exposure, criticality, and controls—so that enterprise teams can organize the best mitigation strategies. Such mitigation can start with those vulnerabilities that can have the greatest negative impact to critical assets.

Ultimately, the goal is to identify a breach likelihood calculation, which is a computed summation of the individual attack vector computations. Such analysis is complemented by probabilistic graph models which estimate the vulnerability levels associated with the various risk scenarios. Collectively, these computations and values provide an organization with an accurate understanding of their security posture.

Cyber Risk Quantification in Dollars

The goal of accurately establishing a quantitative measure of security posture for the organizational attack surface requires use of a risk formula that makes sense to the local domain. To avoid multiple equations, formulas, and other metrics, the Balbix platform defines a consistent cyber risk equation that can be used across all assets and over all aspects of the organization to identify a meaningful posture assessment.

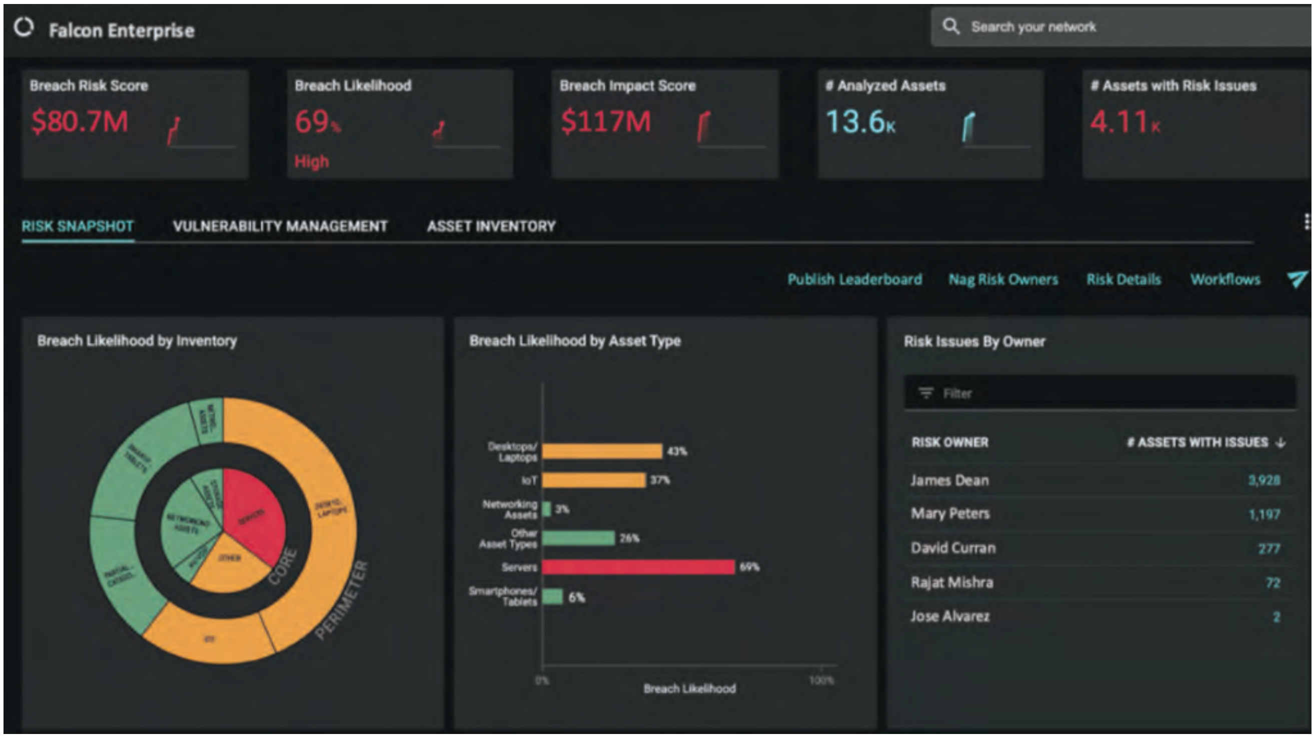

Falcon Enterprise

Search your network

Breach Risk Score: $80.7M

Breach Likelihood: 69%, High

Breach Impact Score: $117M

Analyzed Assets: 13.6K

Assets with Risk Issues: 4.11K

RISK SNAPSHOT | VULNERABILITY MANAGEMENT | ASSET INVENTORY

Publish Leaderboard | Nag Risk Owners | Risk Details | Workflows

Breach Likelihood by Inventory

Breach Likelihood by Asset Type

Risk Issues By Owner

Filter

RISK OWNER	# ASSETS WITH ISSUES ↓
James Dean	3,928
Mary Peters	1,197
David Curran	277
Rajat Mishra	72
Jose Alvarez	2

Figure 4. Balbix Platform—Risk Quantifications

The Balbix platform automates the calculation of risk in dollars. While this is certainly not a new strategy in enterprise cybersecurity, the specialized artificial intelligence models integrated into the platform support the calculation of risk trending, breach likelihood, breach impact scoring, breach likelihood by inventory, and more. These are presented in a visual display that is easy to share with both practitioners and executives.

Cyber Risk Visibility and Board Reporting

The final goal of the Balbix platform is to ensure that enterprise security teams have the best available tools for reporting and explaining vulnerability and risk posture to the organization. This must include reports for senior executives including board members as well as colleagues with more detailed understanding of security programs. Such reporting must cover the entire attack surface and must account for continuing change.

Most executives will tend to focus on the impact of potential breaches, because this represents the most direct consequence of cyber risk to business operations. Balbix supports detailed impact modeling that uses impact estimates based on several factors, such as prior information, contextual impact modeling based on current usage, volumes, and interactions.

ENTERPRISE ACTION PLAN

It is recommended that enterprise teams act immediately to review, address, and improve their cybersecurity posture assessment. This is best done using an automated platform that can unify existing posture-related tools such as scanning and security testing. As suggested above, the Balbix platform provides excellent support in this regard and should be included in source selection plans.

[1] See https://www.balbix.com/.

[2] See https://www.fairinstitute.org/ for information on how the FAIR (Factor Analysis of Information Risk) model supports consequence analysis based on financial impacts.

Balbix®

Balbix enables businesses to reduce cyber risk by automating cybersecurity posture. Our SaaS platform ingests data from security and IT tools to create a unified view of cyber risk in dollars. With Balbix, you can automate asset inventory, vulnerability management and risk quantification, leading to lower cyber risk, improved team productivity and tool cost savings.

TAG CYBER DISTINGUISHED

REPRINTED FROM THE TAG CYBER SECURITY ANNUAL