



The 'Hyperscale SOC' and the Minds Behind It:

A Machine-learning Foundation for Effective Cybersecurity

CONTRIBUTORS:



**Adam
Mansour**
Head of Sales
Engineering



**Alexis
Yelton**
Head of
Data Science



**Brenna
Gibbons**
Data Scientist



**Jennifer
Mitchell**
Head of Operations



**Luke
Wolcott**
Senior Data
Scientist



**Sean
Hittel**
Distinguished
Security Engineer



**Perry
Spector**
Data
Scientist

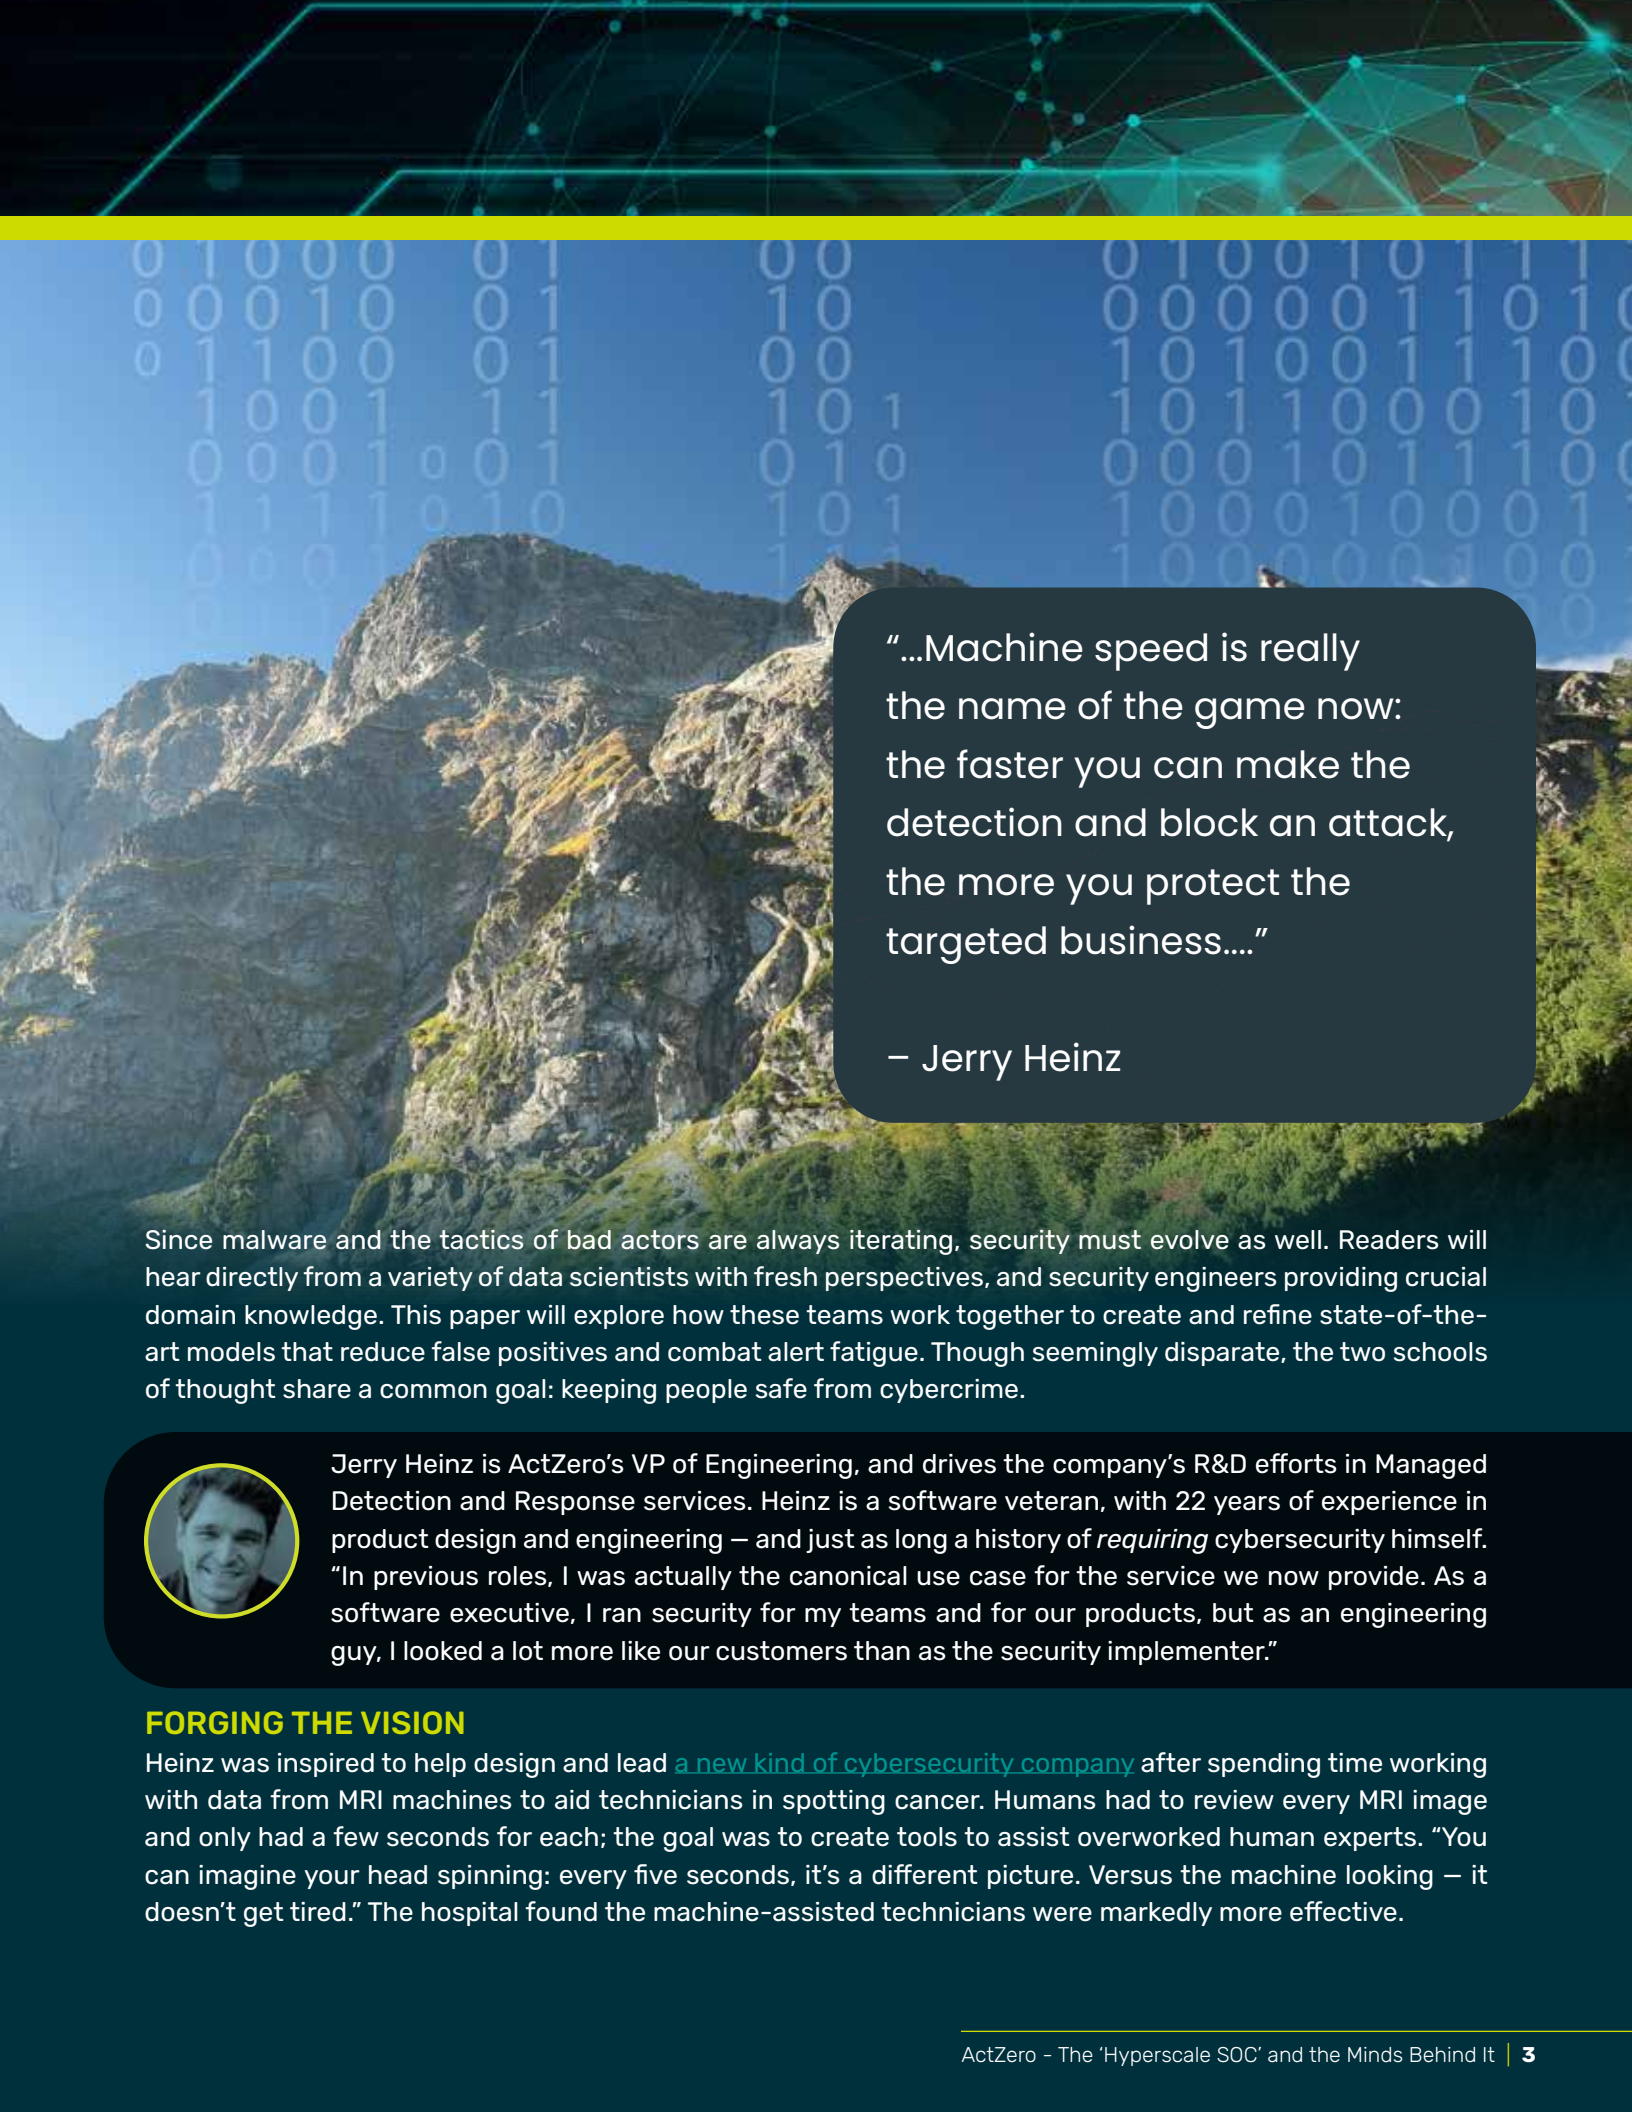


**Jerry
Heinz**
Head of
Engineering

1 Introduction

With a plethora of data, configurations, and potential attack vectors across the modern midsize enterprise, the sheer volume of available information can be overwhelming for those tasked with securing their organizations. Alerts are snowballing, visibility is limited, and even well-equipped organizations are falling victim to advanced threats. Compounding the matter, malicious actors are operating at machine speeds. How can a Security Operations Center (SOC) possibly keep up?

Enter the “**Hyperscale SOC**,” an innovative approach centered on data science and powered by a company that was purpose-built to put data first. This paper will examine how data science empowers human threat hunters, improves the signal-to-noise ratio through high-fidelity alerts, and allows for scalability. Modern threats like ransomware demand the fastest possible reactions, and data science is requisite to respond to attacks at machine speeds. Despite the hype of artificial intelligence, however, data science alone is no panacea – machine learning must be properly implemented and maintained, not just “bolted on” after the fact.



“...Machine speed is really the name of the game now: the faster you can make the detection and block an attack, the more you protect the targeted business...”

– Jerry Heinz

Since malware and the tactics of bad actors are always iterating, security must evolve as well. Readers will hear directly from a variety of data scientists with fresh perspectives, and security engineers providing crucial domain knowledge. This paper will explore how these teams work together to create and refine state-of-the-art models that reduce false positives and combat alert fatigue. Though seemingly disparate, the two schools of thought share a common goal: keeping people safe from cybercrime.



Jerry Heinz is ActZero’s VP of Engineering, and drives the company’s R&D efforts in Managed Detection and Response services. Heinz is a software veteran, with 22 years of experience in product design and engineering – and just as long a history of *requiring* cybersecurity himself. “In previous roles, I was actually the canonical use case for the service we now provide. As a software executive, I ran security for my teams and for our products, but as an engineering guy, I looked a lot more like our customers than as the security implementer.”

FORGING THE VISION

Heinz was inspired to help design and lead [a new kind of cybersecurity company](#) after spending time working with data from MRI machines to aid technicians in spotting cancer. Humans had to review every MRI image and only had a few seconds for each; the goal was to create tools to assist overworked human experts. “You can imagine your head spinning: every five seconds, it’s a different picture. Versus the machine looking – it doesn’t get tired.” The hospital found the machine-assisted technicians were markedly more effective.



2 Security is a data problem: The data-first approach

THE ISSUE WITH 'BOLTING ON' DATA SCIENCE

Heinz realized that threat hunters could also benefit from machine learning in reviewing raw information, so he set out to do the same thing for cybersecurity. He explains that when ActZero was in the conceptual stages, the team chose to break with the industry trend of treating data science as an afterthought. Instead, they would create the new hyperscale SOC by building on a foundation of data science and automation.

"Security is a data problem," says Heinz. "Cybersecurity companies have tried to bolt on data science after the fact – they all see this is where the attackers are going, and that is what's necessary to respond quickly to changes in the environment. Machine speed is really the name of the game now: the faster you can make the detection and block an attack, the more you protect the targeted business from suffering any financial or reputational harm."

THE CHALLENGE: RETROFITTING MACHINE-LEARNING CAPABILITIES

However, this shifting approach to security is not something that can be simply riveted on retroactively, says Heinz. "We saw many security companies embracing a so-called 'AI' approach, and that everyone wanted to say 'Oh, I'm going to use data science for this.'" But data science is a field that requires expert researchers with years of training. "This is not some algorithm that you can just pull off the wall and plug in. It requires a structured way of thinking and a proper foundation."

With this approach in mind, Heinz and his team decided to do something different: form a new company *starting* with data scientists, and then build engineering and R&D teams around their vision. (More on this later in the "Fresh perspectives in cybersecurity" section.)

ACQUIRING HIGH-QUALITY, DIVERSE DATA

With this plan for a solid data-driven infrastructure and a leadership team in place, the next step was to find a company that had been successful under the increasingly outdated security-first model. "We were looking for a company that had the customers, the security knowledge and expertise, the right culture fit, the right data diversity," says Heinz. After much searching, the team [found IntelliGO Networks](#).

IntelliGO's high-quality data took years to collect. "Something that IntelliGO got very right was pulling in a large swath of data," says Heinz. "They had many different kinds of customers and log sources, so they also had excellent data diversity. You need that in order to produce a robust machine-learning model. Otherwise, it will be weighted one way or the other, skewed towards a specific type of environment, threat, or organization."

IntelliGO also came with a major hurdle to overcome in terms of the operation's scalability, familiar to IT leaders attempting to build their own SOC: as the company grew, their security professionals were struggling to keep up with the ceaseless torrent of data, logs, and alerts. "Frankly, they were overwhelmed with the amount of work. That was expected – that's what data science could help with."

3 The push to scale security operations

Adam Mansour is ActZero's Head of Sales Engineering, with over 15 years of experience in the cybersecurity sector. Mansour founded IntelliGO Networks in 2005, where he developed a proprietary MDR platform to enable threat hunters and served as the company's CTO prior to its acquisition by ActZero. Mansour has helped build countless SOCs, and knows firsthand the pitfalls of attempting to "bolt on" machine-learning capabilities after the fact.



THE CHALLENGE: ALERT FATIGUE

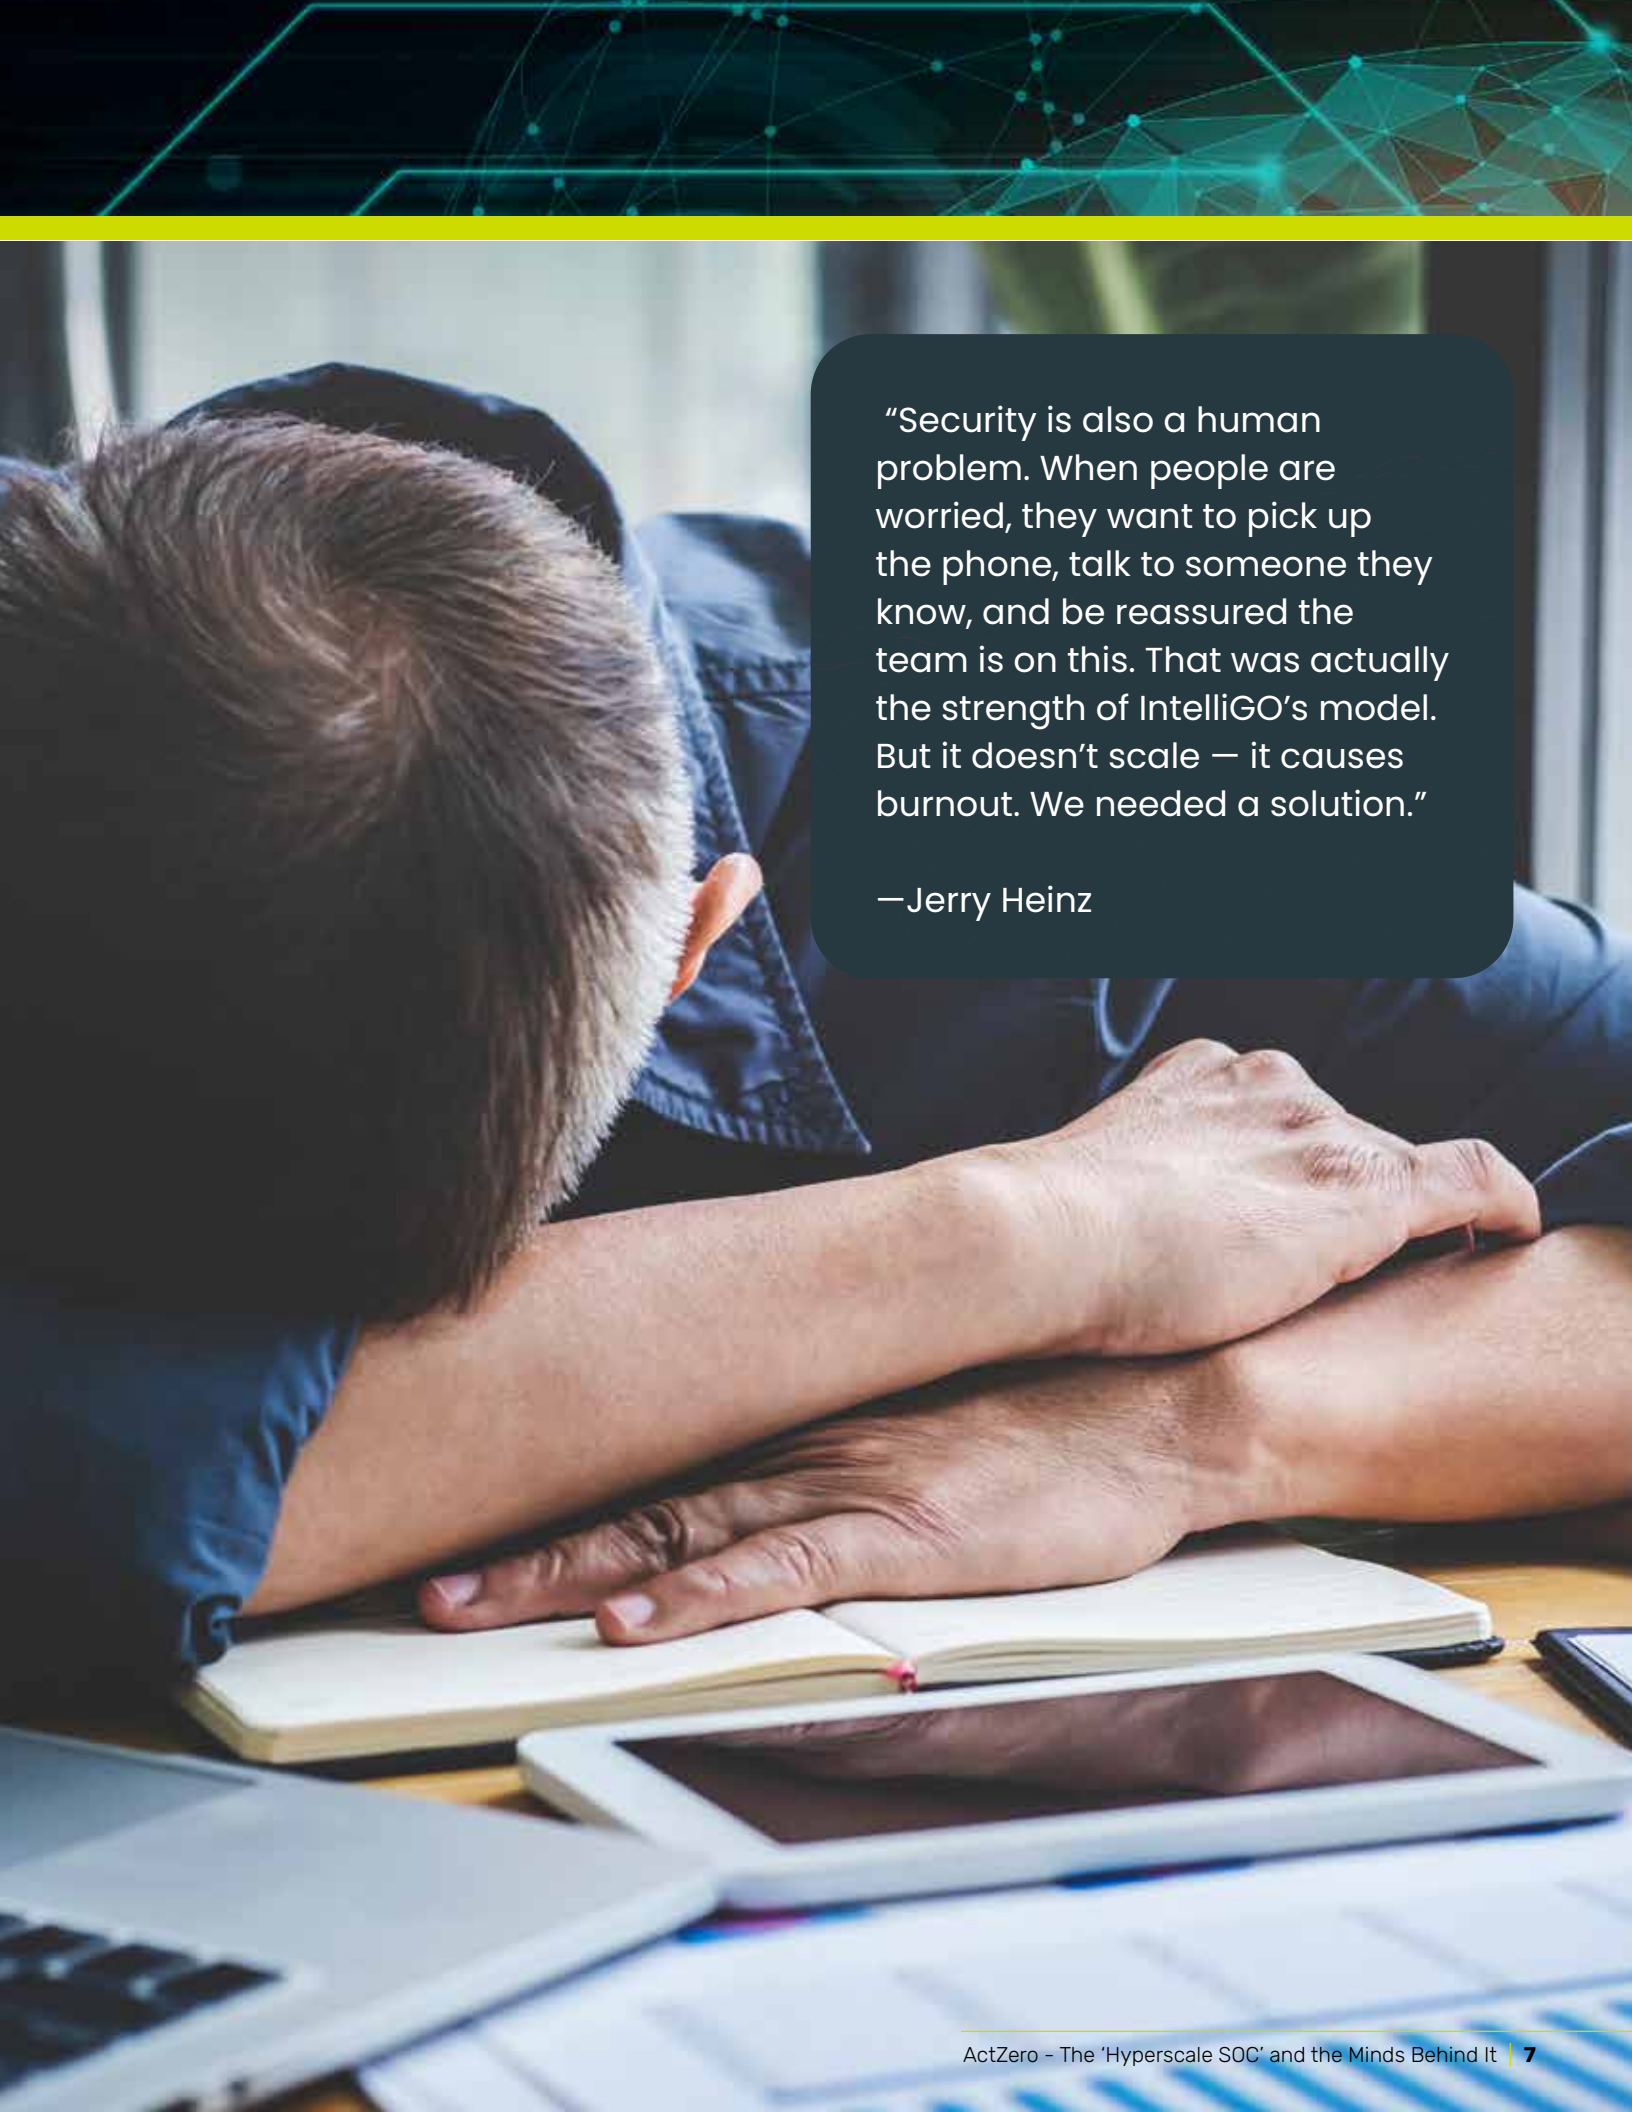
Mansour says that as a security company grows, so does the risk of alert fatigue – a challenge many IT stakeholders are all too familiar with. Your security operation may receive hundreds, thousands, or even more alerts a day – “and each analyst is able to properly handle 10 of them. With this influx of alerts, relevant stuff can easily get missed. You're holding on by a thread, wondering what you can do to alleviate the stress and reduce alerts so you can get to a better state.”

Like Heinz, Mansour realized the potential of introducing machine-learning capabilities to aid the beleaguered humans in the SOC – and the hazards of doing so after the fact. “When you go to bolt on this technology, you fall into a trap. Without consummate data professionals on your team, you end up with a machine-learning algorithm that you can't modify, which means your protections don't get better over time,” says Mansour.

“The fundamental flaw in *adding* machine learning to your stack is that you don't redesign your SOC or the roles of the humans working inside it so the models can keep improving. Threat hunters need to investigate and action alerts surfaced by the machine-learning detection, but also need to look for anything the models may be missing. And if your security engineers don't test the model full time, it won't improve on its own from the raw data set.”

Mansour points to another issue with implementing off-the-shelf machine learning tools. “If you simply leverage whatever product a vendor is giving you for machine learning, you are not participating in the data science – the data goes directly to an analyst. There are no data scientists or security engineers in the loop; there's only an analyst receiving a machine-learning detection without the context from your organization.” Mansour says this type of out-of-the-box model does little to bring down the rate of false positives, and so ultimately isn't much help in dealing with scalability and alert fatigue.

The only way to build a fully capable, machine-enabled security operation, Mansour realized, was to start from scratch. Only this time, bringing in data scientists from the beginning to design custom models, rather than trying to plug in someone else's.



“Security is also a human problem. When people are worried, they want to pick up the phone, talk to someone they know, and be reassured the team is on this. That was actually the strength of IntelliGO’s model. But it doesn’t scale — it causes burnout. We needed a solution.”

—Jerry Heinz

4 The data scientist revolution: the minds behind the machine

Alexis Yelton is ActZero's Head of Data Science, where she drives the development and training of machine-learning detection models. She holds a PhD from the University of California, Berkeley in environmental science with a focus in bioinformatics, and carried out postdoctoral studies at MIT. She has helped bring data science to various industries, but this was her first foray into cybersecurity.



THE CHALLENGE: EXPANDING THE ROLE OF MACHINE LEARNING IN SECURITY

When she was first offered the role, Yelton says “they pitched it to me ‘as cybersecurity is really far behind some of the other technical fields in making use of machine learning. We want to bring in data scientists who don’t have a security background to modernize the field.’”

Yelton seized the chance to help architect a data-driven security company. By designing and incorporating models that use large amounts of high-quality, diverse data to establish a baseline, Yelton is helping security professionals in the hyperscale SOC determine exactly what “normal” activity looks like. The models are trained to detect unusual activity and bring it to the attention of security analysts.

The data science team takes a multi-pronged approach to threats, building **anomaly detection** and **supervised models**, as well as **holistic models** across multiple data sources (such as endpoint,



network, and cloud) to better cover the attack surface. As she explains, “With anomaly detection, you gain the ability to detect new attacks and attack families, and with supervised models, you get a strong signal as to what existing attack types look like. Holistic models add to these well understood techniques to ensure that weaker signals are acted on where they would otherwise be missed.”

The team also leverages what they call “in-environment models”. These are models specific to certain customers, to bridge the gap from general to specific. By factoring in data across the entire customer base, they get the best of both worlds - detections specific to that customer, but trained by the whole data set. Without this in-environment tailoring, a model would have to permit some known false negatives (missed attacks) so that it does not false positive. False positives not only contribute to alert fatigue, but prevent automation of responses, as they can disrupt business operations.

SIGNAL-TO-NOISE RATIO

The security status quo has been sending more alerts than humans can feasibly handle, as any reader who has used a SIEM (managed or otherwise) can attest. The goal is high-fidelity detections that are truly indicative of malicious behavior, thus empowering threat hunters to find all the attacks without a lot of false positives. This improved signal-to-noise ratio helps tremendously with the issue of alert fatigue, says Yelton. “When you’re going through that many alerts, you’ll miss alerts that are actual attacks.” (For more on signal-to-noise indicators, check out our white paper “)

It’s not just security companies dealing with scalability that risk missing signal. “In fact, there have been some high-profile breaches in which big enterprises lost hundreds of millions of dollars from attacks that were actually detected, but the admins or analysts looking at it did nothing. We reduce those false positives to prevent that from happening, without sacrificing signal.”

With the foundation of sound data, a multi-pronged approach, and a variety of models, and a driving principle to enable threat hunters, operationalizing the SOC would be critical.



5 Operationalize it: Efficacy and Efficiency

As Head of Operations, Jennifer Mitchell drives ActZero's strategic efficiency and scalability, and specifically oversees the management and architecture of the SOC. Coming from IntelliGO, she was all too familiar with the issues surrounding the old model. But, with new capabilities afforded by data science, new avenues were at her disposal to create a truly effective and efficient security operation.



“The hyperscale SOC involves the orchestration of numerous components. Automations, high-efficiency workflows and structures, machine learning detections, innovative operationalization and workplace psychology all play a role. Machine learning is certainly critical, but the whole is so much more than the sum of its parts,” said Mitchell, for whom making the hyperscale SOC a reality remains her primary responsibility.

THE CHALLENGE: REIMAGINING PEOPLE, PROCESSES AND TECHNOLOGY IN THE SOC

Mitchell had the foresight not to pursue the 'bolt-on' approach. "We didn't 'just re-engineer' it either," she said. "To operationalize for hyper-scalability we had to completely deconstruct security operations, and build them back up, from scratch. A complete reimagination! This included creating new processes to achieve aggressive hyperscale-specific goals on both accuracy and efficiency." Yet, the scope of the 'ground-up' alternative was vast.

DATA-FIRST, OUTCOME-FOCUSED ASSESSMENT

Mitchell leveraged operations first-principles, often overlooked at technology companies. "We did a full standardisation and value assessment on SOC activities, ensuring we prioritised time spent on high-impact security activities." She crafted innovative assessment methods. "We gave each person in the SOC physical devices for tracking activities - imagine dice, but with eight sides, and each side labelled with a particular type of task," she said. "It allowed us to be granular with activity data, and gave threat hunters an easy way to track. Plus, participation was extremely high - having it on your desk is a constant visual reminder to switch."

The output allowed Mitchell to identify the best opportunities for automation; the places where automation stood to save SOC personnel the greatest amount of time, without sacrificing quality of response. "The benefit of having the right instrumentation was requisite for getting the most out of the machine-learning detections," explains Mitchell. "The feedback loop it affords really integrates the security stakeholders and data scientists. You don't just want that intelligence to [drive higher quality alerts](#), but also efficiency - you need the systems and feedback loops to operationalize the intelligence across the entire SOC. That's **what makes it hyperscale**. The misconception is that you can manage it the same way as a traditional SOC."

As such, Mitchell's management practices and the psychology of employees played an increasingly relevant role in integrating all the elements. "The threat hunters lived in an ever-evolving threat landscape, which meant their job functions often changed. I needed them to be effective change agents in a space that is not normally known for this. So, we also gathered data on the effects that context switching, specialisation, randomisation, and even the time of day can have on the productivity and efficiency of the SOC," she said. "You would be amazed at the changes in analyst's satisfaction that we now attribute to the resulting changes."

SYNERGISTIC APPLICATIONS

There was also an ongoing identification of automation use cases to reduce manual work wherever possible. These included both traditional process refinement and software automations, as well as new ones afforded by machine-learning detections. "With so few false positives, we were able to automate responses to high-severity, high-fidelity indicators of attack [IOAs]," said Mitchell.

Mitchell highlights the virtuous cycle this reimagination brings about. "The hyperscale SOC is never actually finished. It's a continuous wheel of assessments, collaboration, improvements, and orchestration. We needed the accuracy to enable the automation; we needed the automation to enable the efficiency; the efficiency yielded a depth of focus, to tackle threats that would have been missed in a SOC subject to alert fatigue. How I manage these factors to maintain the right balance is the art. Anything less than all of them is insufficient - you could be the most efficient at detecting only one thing - and that's woefully inadequate. By doing this all together, we're able to do this at a meaningful scale, hence 'hyperscale.'"



6 The PhDs in the SOC

Building a hyperscale SOC that actually produces high-fidelity detections takes great minds — and equally great data. Here are some areas in cybersecurity where specialists are continually improving the signal-to-noise ratio. Read on for how they empower threat hunters and security engineers, and more details on their approaches.



Luke Wolcott is a Senior Data Scientist at ActZero, with a PhD in mathematics. He was excited to join the data team when he realized the cutting-edge nature of the work to combat ransomware. “That’s what you do as an academic,” says Wolcott. “You take on new areas, you do the literature review, and then you say, ‘Okay, I’m at the frontier now — how can I push the frontier?’”

MACHINE-SPEED RANSOMWARE DETECTIONS

One of Wolcott’s current projects is harnessing what he calls “the magic of machine learning” to find subtle patterns in ransomware attack data. By identifying a variety of “quiet little signals” and assigning each a weight based on statistics, he works with ActZero’s security engineers to help the models learn what is a true early indicator of a ransomware attack — and thus when to take immediate action. “If a certain detection goes off, we definitely want to quarantine instantly — it’s a question of severity,” he says. (For more on how the ability to respond at machine speeds can dramatically reduce the damage done by ransomware attacks, see our white paper [“The Rise of Ransomware-as-a-Service.”](#))



IDENTIFYING SUSPICIOUS SCRIPTING

Brenna Gibbons has a PhD in materials science, and is one of the data experts working behind the scenes in the data-driven SOC. She is presently focused on aiding computer models to better identify signs of suspicious scripting – without generating too many false positives. The models have access to data about thousands of threats, gathered from both public databases and customers. By comparing this to the data found in presumably clean environments, the machine-learning models are able to identify new potential indicators of attack.



Gibbons contributes to detecting PowerShell and command-line attacks. Such assaults can be devastating, she says, but are also relatively rare – something of a mixed blessing for a data scientist working in security, for whom attack data is integral to the success of models. She worked alongside a security team that generated “tens of thousands of legitimate attacks using PowerShell and other scripting tools,” in order to produce more usable data. (For specific examples, [see our blog post here.](#))

DEFENDING THE CLOUD

Perry Spector is another data-science expert, focused on cloud detections. Like his peers, Perry’s academic research was in a markedly different area. “I studied earth science and the ice sheet that covers Antarctica. Right now, I’m working on identifying malicious attempts to log into Microsoft 365 accounts and looking for unusual patterns in user behavior,” he says.



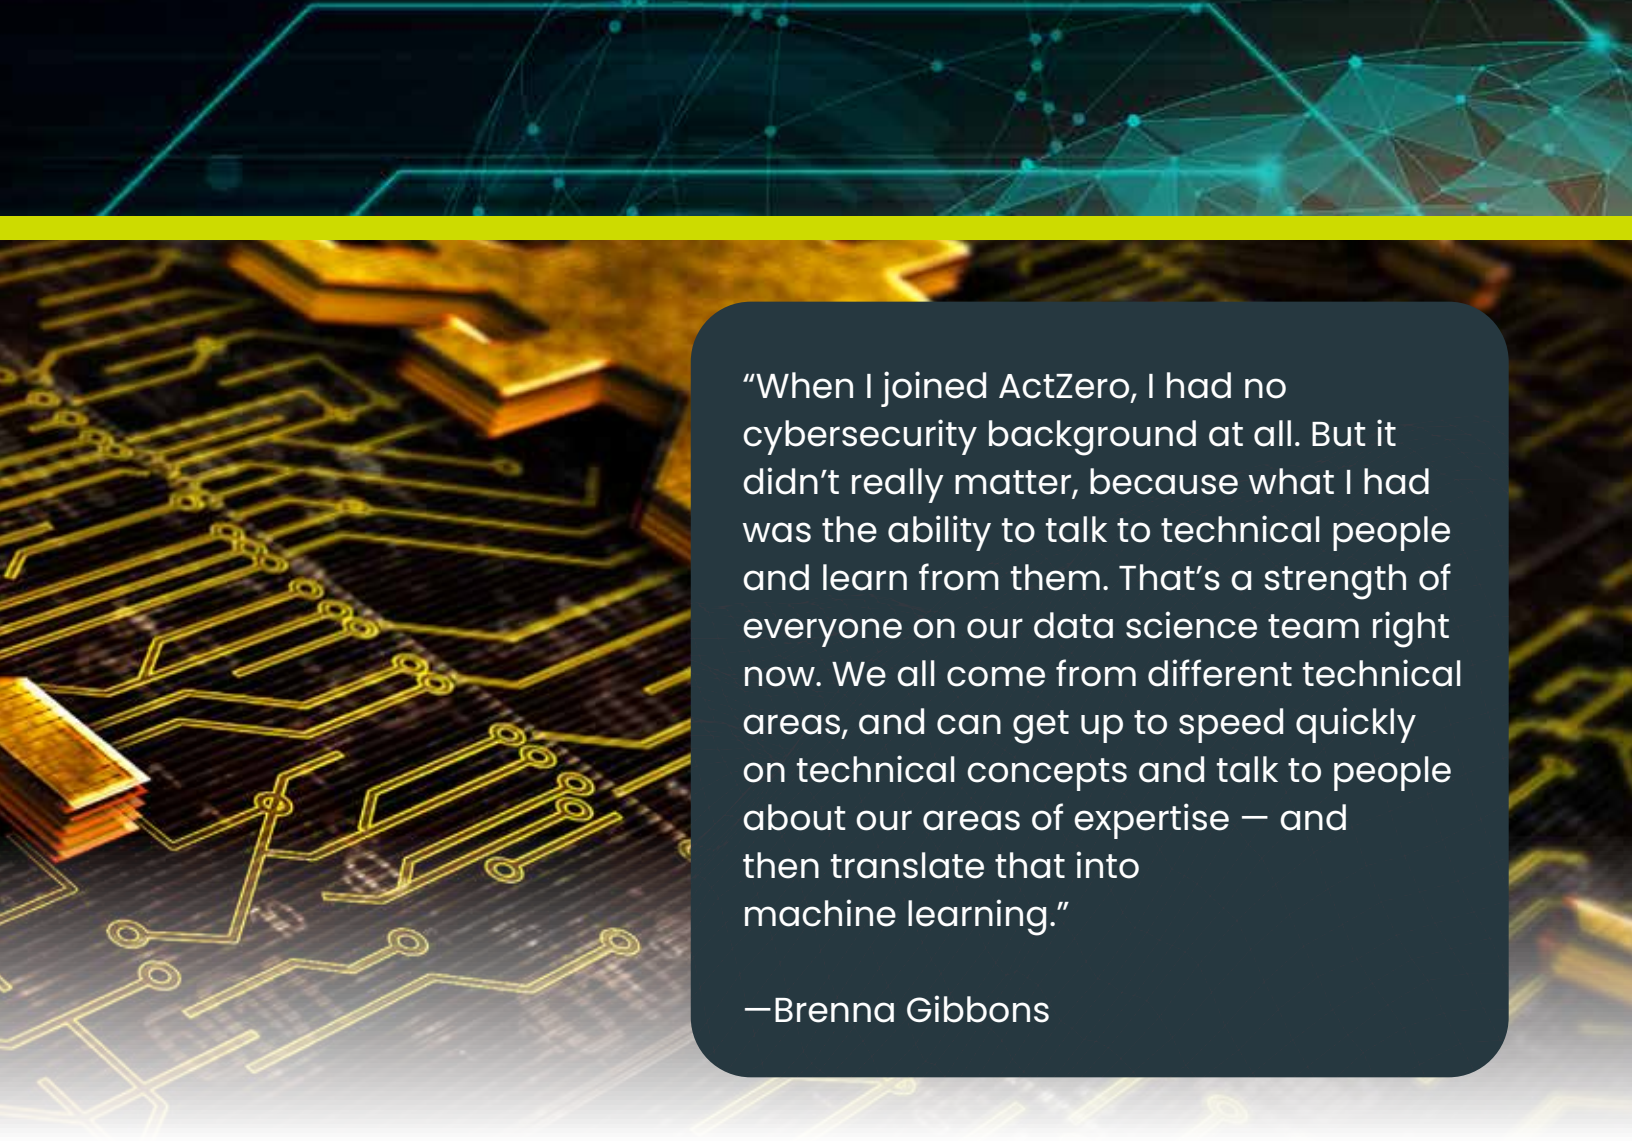
One challenge for Spector is using machine learning to identify truly anomalous user behavior in the Microsoft cloud. “For example, take geolocation in login data,” he explains. “You’d think it would be suspicious if someone logged in to an account from somewhere halfway around the world – but it could just be that the user is on a VPN or on vacation.” Again, the goal here is a better signal-to-noise ratio. “Part of how I use data is to find and analyze genuine attacks, and then randomly sample them to build up a synthetic data set that can be used to train a model to classify logs as either benign or suspicious.” (For more on this, see our white paper “[Securing the Microsoft Cloud from Azure to M365.](#)”)



MOVING BEYOND MALWARE

Cybersecurity had already made strides in using data to comprehend malware through static or dynamic analysis, says Yelton – but with modern attacks, malware-focused prevention technology is no longer sufficient. “Malware is just a small part of the threats that we see. Often, the attack sequence involves many components that can be accomplished without relying on malware. For example, [‘Living Off the Land’ attacks](#) rely on systems often already present in environments to do reconnaissance, steal credentials, move laterally, and establish persistence.”

Yelton says the reason existing malware-detection models are inadequate is due in part to the lack of data science perspectives in their creation. “A lot of those models were built by security engineers, which I think is great, but they are subject to confirmation bias; they’re looking for what they already know. This is more obvious to people who don’t have a security background, and can look at the data with fresh eyes.” Her team is using a data-driven approach to get the whole picture around attacks and arrive at new solutions to old problems.



“When I joined ActZero, I had no cybersecurity background at all. But it didn’t really matter, because what I had was the ability to talk to technical people and learn from them. That’s a strength of everyone on our data science team right now. We all come from different technical areas, and can get up to speed quickly on technical concepts and talk to people about our areas of expertise – and then translate that into machine learning.”

—Brenna Gibbons

FRESH PERSPECTIVES IN CYBERSECURITY

Gibbons also argues for the sense of bringing data scientists into the world of security, pointing to the well-documented shortage of cyber talent. “There are already too few threat hunters and too few security experts. So turning them into data scientists isn’t necessarily the best use of resources.” (For more on this, [see our blog entry here.](#))

She draws a parallel to her own work. “If you’re trying to train a machine-learning algorithm to solve some particular problem – for example, with neural networks – the thing that makes them so powerful is that you don’t tell them exactly how to solve the problem. You just give them data, and they sort out what the best features are as part of training. In some ways, that’s what we’re doing. It’s human learning, but it really echoes machine learning in that we’re letting the data drive our exploration. Someone with a deep cybersecurity background might have very set ideas about how to solve a particular problem and make the data fit that. We’re coming at it from a different direction,” says Gibbons.

No one would agree with that more than Mansour. He credits Yelton with helping make a variety of decisions at the outset that wouldn’t have occurred to traditional security professionals. “Alexis spent weeks and weeks flying me back and forth to Seattle to explain that she didn’t have enough data yet for the models,” recalls Mansour. “But we had terabytes of data! What she meant was that we didn’t yet have all the right *attributes* within that data to help the machine make the decisions we needed it to make. Our data practitioners still had to augment our data streams to get the algorithms to work properly. Through their custodianship of the data, they could tell me what information was relevant to security, rather than me simply telling them, ‘Here are the valuable points that we use to make a decision.’”



THE CRITICAL IMPORTANCE OF DATA

Wolcott says the most important thing to data scientists is that working with good data they find interesting. “If you don’t have data, then you’re just wasting your time.” During his job interview, the first thing he asked was whether ActZero had the quality, structured data he needed to succeed. “By acquiring IntelliGO and their customers, ActZero automatically had ample data. Once I heard that it was like ‘absolutely, I’ll take the job – you actually have data, you have a lot of it.’ Having good-quality data and having a lot of it is just something that data scientists fall in love with.”

THE CHALLENGE: GETTING THE RIGHT KIND OF DATA

Yelton argues that most security companies serving smaller enterprises don’t have the right kind of data, or enough of it. With access to a trove of data and security experts to train and validate the models, this was a problem Yelton was eager to tackle. “There’s a ton of data out there, but the issue is there’s not a lot of *threat* data. You need the actual data of threat actors attacking your customers, and you need a lot of it.” (She talks more about how ActZero approaches gathering this data in her appearance on [the Data Standard podcast](#).) The right kind of data meant both a high volume and highly variable log data, across multiple sources and vectors – endpoints, networks, the cloud, and security technologies.

One wrinkle she helped iron out was the case of smaller companies that have never been breached. What did the models do without data? “What we *do* have is a lot of data on ‘not breaches,’ and that’s why we don’t only build models that look specifically for threats,” says Yelton. “We build those supervised models, but we also build unsupervised models that just try to understand your data overall, and look for anomalies. So, when we’re dealing with [such companies], we still have a large amount of information about what’s normal for them. And we’re therefore able to confidently determine what’s abnormal.”

Armed with excellent data, the PhDs in the hyperscale SOC must work closely with security experts, who provide domain knowledge and guide the data scientists’ efforts. Here’s how the two seemingly disparate camps come together.

7 The bridge between security and data science

Gibbons recalls her background doing experimentation in material sciences. As her interest in data grew, she realized the size of the gap between the domain experts – the chemists doing the actual experiments – and the computer scientists working on machine learning. “There was so much room for relatively basic data science and machine learning to make a big impact in these fields, but they didn’t know how to talk to each other *at all*,” she says. The people who understood machine learning didn’t know anything about chemistry, “so they’d get obvious things wrong. And those who knew the chemistry just didn’t have the time or the inclination to learn the data science.”

CONNECTING DISPARATE FIELDS

It’s easy to imagine the same issues arising in the hyperscale SOC. Enter Sean Hittel, the Distinguished Security Engineer at ActZero who helps bridge the gap between domain experts and data scientists by living in both worlds. Hittel has been incorporating machine learning into threat-detection engines since the late ‘90s, a mission that has often proved frustrating.



“The bane of my career has been the thinking of people that are pure security – who see me as a data-science zealot, always trying to shoehorn in machine learning,” says Hittel. “And then the data science people think I’m a security fanatic, always trying to solutioneer. The two camps are pretty far apart, to be frank – sometimes ideas that are otherwise good just die because of the fighting.”

THE CHALLENGE: FACILITATING COOPERATION

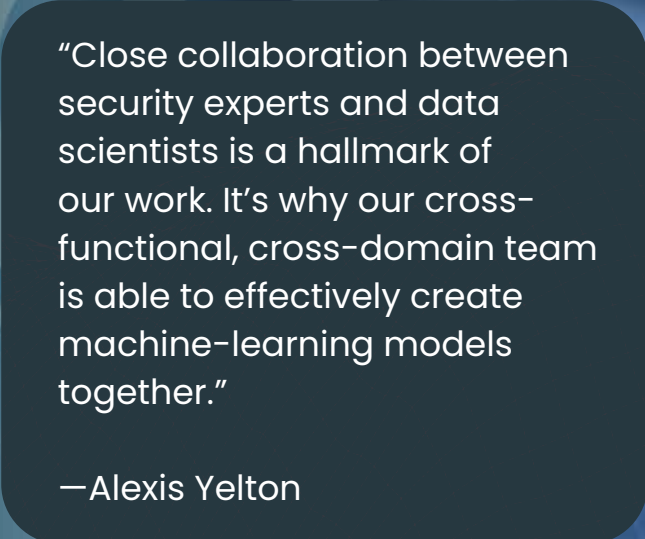
In but one of many functions, Hittel serves as a liaison between the threat hunters and the data scientists, making sure crisp data science is backing up strong security understanding and vice versa. “We spend a lot of time together talking about the security areas that off-the-shelf – or even state-of-the-art – detection engines just can’t do well,” he says. “Then we collaboratively figure out the most important areas to invest in, and determine the best approaches.”

He cites efforts to identify long-cycle attacks as an example. “We know that currently, things like Living off the Land attacks are poorly served by automated mechanisms. And it’s not a matter of building a better mousetrap – attackers know to sneak around that mousetrap. So the security people and data scientists ideate novel approaches to catch attackers, not just attacks, so they can’t simply change their tactics.”

Hittel works with the threat hunters to identify anything the models may have missed, feeding this back to the scientists, thus [refining the efficacy of the SOC](#). Looking at threat hunters’ investigation-and-response processes also yields opportunities for the machine to learn more effectively by providing a closed-loop system of labels that indicate every time a model was right or wrong. It also allows the data scientists to identify opportunities to automate responses the SOC would normally take as a matter of course. Through this close cooperation, the teams have come to appreciate that while they may differ in tactics, they share a common goal: protecting people by stopping cyber threats.

Of course, the data and security experts can disagree, or even butt heads at times – but that was the point from the start. “We want a bit of an adversarial relationship when solving problems,” says Hittel. “If you create a competitive environment, you’ll often end up with a better solution than you would have with consensus.”

Yelton says this slightly antagonistic approach “definitely created some tension at first. But we came to understand just how many really useful things the security experts had for us to consider while building models. And they’ve learned we’re able to find the additional signal in the data that has nothing to do with what they understood to be important or useful. This really helped us come together.”



“Close collaboration between security experts and data scientists is a hallmark of our work. It’s why our cross-functional, cross-domain team is able to effectively create machine-learning models together.”

—Alexis Yelton



8 Conclusion

The need to respond at machine speeds is undeniable in today's cybersecurity landscape. Alert fatigue and talent shortages are serious issues facing modern SOC's, and the failure of older security models is evinced by dire news headlines. You have heard firsthand why common efforts to "bolt on" machine-learning capabilities after the fact are inadequate – and why a collaborative approach that prioritizes data is essential to ensuring high-fidelity detections that evolve with time. You've also heard the value of fresh perspectives on old problems, and how ActZero's data-driven, machine-enabled, **hyperscale SOC differs where it matters.**

Just as providers can't simply bolt on cutting-edge data science, it's prohibitively difficult to do this in-house. Acquiring and managing the high-quality, diverse data needed for successful models is a major undertaking outside the scope of midsize organizations (not to mention the people to undertake this). However, IT stakeholders like you can still take action.

Ask your providers about signal-to-noise, and how they're taking a data-first approach. Ask your security stakeholders what data sources they're utilizing. Are they responding to every single alert? Do they have enough confidence in the fidelity of their detections to enable automated responses?

Old security solutions buckled under the weight of too much data; [ActZero's approach](#) only gets stronger with more data. As Heinz says, there is a symbiotic element to the data-first approach that means everyone benefits: "We need customers to provide us data. With great data, we can provide better detections of attacks, which means we can better protect our customers, which then leads to more customers, which then leads to more data."

To see the output of this rigorous process in action, [request a demo of our intelligent Managed Detection and Response service](#)

Or to learn more, check out our [blog](#) and other [resources](#) about how we use machine learning to [identify suspicious scripting](#) [defend the Microsoft cloud](#), and [detect and respond to ransomware](#)



TORONTO

207 Queens Quay, Suite 820
Toronto, Ontario M5J 1A7

MENLO PARK

2882 Sand Hill Road, Suite 115
Menlo Park, California 94025

SEATTLE

925 4th Ave., 20th Floor
Seattle, Washington 98104

