CYBER RESILIENCE

ThinkTank

Sponsored by mimecast

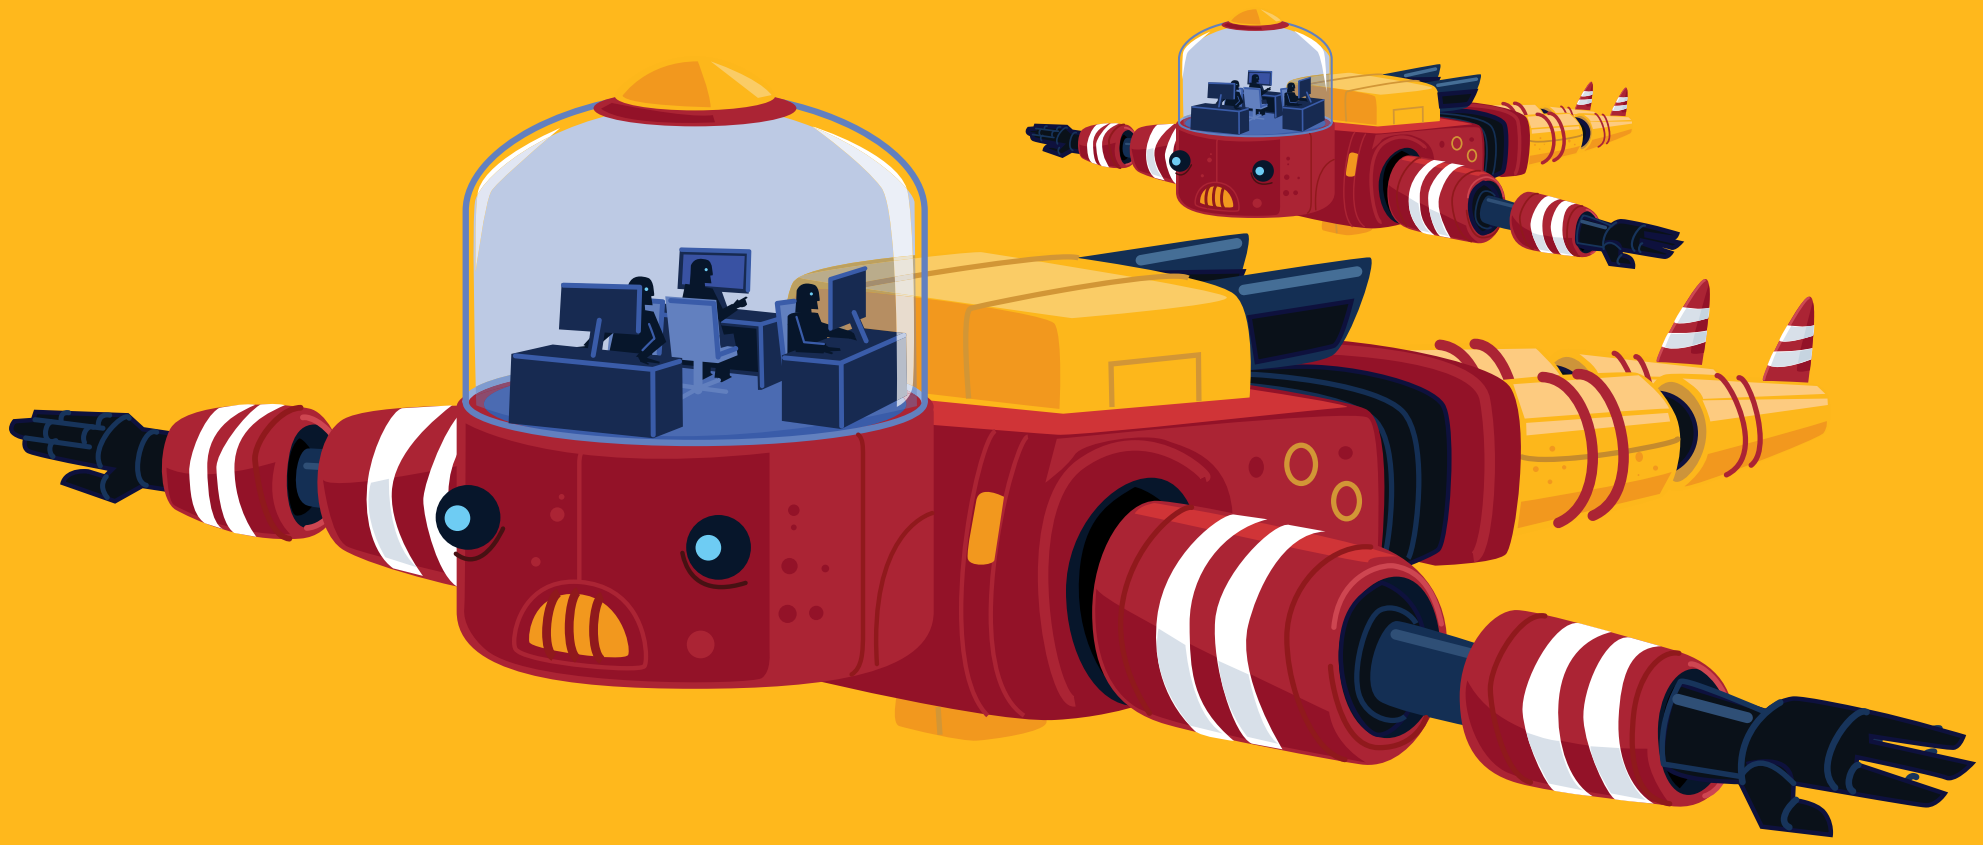

Recommendations

for managing AI risks

01/05

Vendor AI capabilities and automation must be integrated

For CISOs using vendors that employ AI, it's less about what intelligence their tools provide, and more about how the organization's internal security team can integrate external AI tool findings into their broader security programs.

Claus Tepper, head of security operations at South Africa-based financial services company, ABSA group, told peers his team uses trusted vendors to provide detection services across a massive volume of data. But, Tepper's team must have a "very high" confidence level in the AI-gathered intelligence in order to act.

"If we move too quickly on AI and automation and ultimately cause impact to the organization," Tepper said, "I'm going to have to take 10 steps back on my security program."

Other security experts expressed a similar and delicate balance when considering vendor-provided AI attack detections. For HealthFirst CISO Brian Miller, that means using vendors to send alerts to his security information and event management (SIEM) and security orchestration, automaton and response (SOAR) systems.

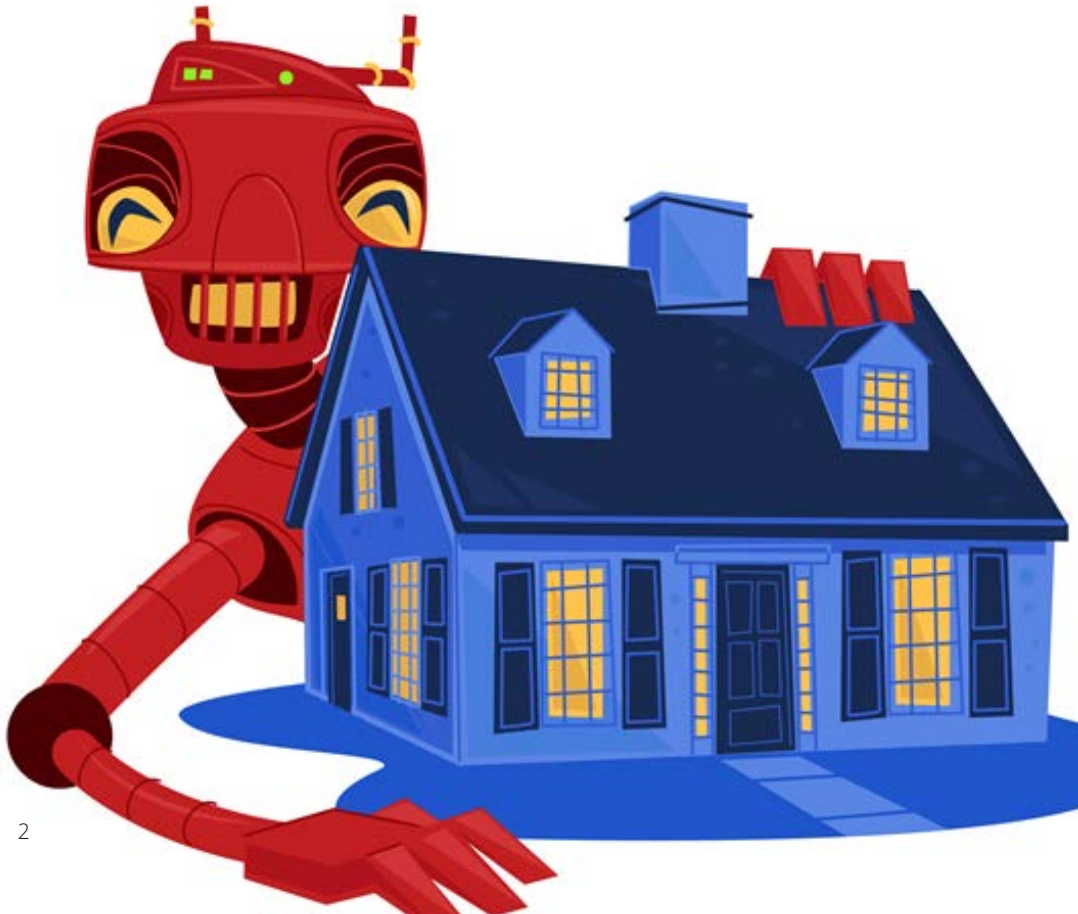

"If we move too quickly on AI and automation...I'm going to have to take 10 steps back in my security program."

Claus Tepper, Head of Security at the ABSA Group

"It's really 'can I make a bet on the vendor?'" said Brian Miller, CISO at HealthFirst. "I rely on vendor products for edge cases, but overall, I'm using them to send alerts to my SIEM/SOAR."

Those alerts raise a flag with Miller's team to decide on an action regarding the AI vendor's attack intelligence.

Key Steps to Managing AI Vendors

1. **Determine the right combination** of AI for your business.
2. **Find the process** for determining a trusted vendor.
3. Ensure that your **AI technology is integrated** with your systems.
4. **Use the combined power of these systems** to identify choke points — like active directory or email — in the cyberattack chain.

02/05

Use AI Technology that Can Take on the Complexities of Human Behavior

Machine learning, a subset of AI, proves effective in detecting highly directed attacks that are difficult for traditional, rule-based systems to detect — and that's because it would be near impossible to create enough/right rules to detect more sophisticated cybercriminal activity, said Paul Everton, Mimecast security architect and founder of the AI-based company, MessageControl, which was acquired by Mimecast in 2020.

Creating algorithms around human behavior is also difficult, but it can be done in a comprehensive way. Mimecast's CyberGraph, for example, compiles all end user touch points and uses those user data points to train machine learning models.

"Seventy percent of the time is used to compile data, 30% is used for building models to get accuracy," Everton said.

Many believe the sheer fathoms of data in today's work environment — see CRMs and cloud-based productivity platforms like Microsoft 365 — make it impossible for security teams to continue to do their jobs without the assistance of machine learning algorithms.

"But machine learning is not a silver bullet," said Peter Tran, Head of Global Enterprise Cyber & Product Security Solutions at InferSight.

70%
of time is used to compile data

30%
of time is used for building models to get accuracy

“I believe there’s intelligence in security. It’s just carbon-based, it’s not silicon based.”

Sam Curry, Chief Security Officer at Cybereason

“I think there needs to be a paradigm shift,” Tran said. The data needs structure, he said, and some logic needs to be added to make an algorithm behave more like a human user.

And it’s important to note the distinction between AI as a field, and machine learning as a *type* of AI.

“I believe machine learning is becoming more mature in some of the things it can do as functions,” said Sam Curry, chief security officer at Cybereason, who’s been a developer and a CTO in past lives, and has worked with AI and machine learning since 2005. “Especially when they’re discreet functions, when there’s few variables, large feedback loops, as opposed to AI. By the way, I don’t know if there’s anyone who ends a statement with ‘AI.’”

That’s to say, AI shouldn’t be used as the catch-all solution to a company’s security issues. Curry is fine with someone saying, “AI, let me explain…,” and he’ll double-click on that a few times. But when vendors say, “Because, AI,” it’s too much like “Magic Bucket,” he said.

“I believe there’s intelligence in security,” Curry said, “It’s just carbon-based, it’s not silicon based.”

“And if we had an AI, who knows if it would want this job?”

Sam Curry, Chief Security Officer at Cybereason

The intelligence is there, it’s undeniably good, but the problem is scale.

“We haven’t solved the Turing challenge,” Curry said. “I’d like to see the carbon intelligence assisted by the silicon intelligence.”

The Turing challenge, otherwise known as the Turing test or “The Imitation Game,” refers to a test designed by computing pioneer Alan Turing in 1950 that established a methodology to determine whether a computer could be considered “intelligent.”

“And if we had an AI, who knows if it would want this job?”

03/05

Look at AI Relative to your Security Strategy and Business Risk

Assuming the algorithm accepts the job (smirk), what level of actions are CISOs comfortable with being performed by machine learning automation?

"I'm more worried about the firewall, because I might shut down payments where we're paying the hospitals, who if they don't get their payment, they can't pay their people," Miller said. "It really is, 'what's the risk to the business?' P.C. (personal computer) is easy, I kick someone off, and it might hurt our quality score a little bit. With firewall I could shut down our business, and you could divide the daily cost of $11 billion — and that's how much business I've shut down."

Marc French, managing director and CISO at Product Security Group, called that "Denial by AI."

"At what point in time does your business shut down because AI made a poor decision?" French said. "Whether that's your customer shutting off your claims processing or shutting your user training data that you're putting in based on tags."

But some consider that as worth the risk, and there are levels.

"At what point in time does your business shut down because AI made a poor decision?"

Marc French, Managing Director and CISO at Product Security Group

"If you shut down my entire enterprise for three days of the week, but if you stop a ransomware attack, I can go back to the board and say, 'this was the right thing to do,'" Miller said. "That's a little extreme. But I could stand behind that."

It really comes down to — are you able to identify the critical nature of the attack from the information the machine learning tool provides? If ransomware is moving laterally throughout an organization's system and the security team disrupts business, that's not going to be a problem, Miller said.

"If someone is doing a Golden Ticket attack or a Kerberos attack on an active directory, I know that that's bad," Miller said. If the AI provides a known indicator of an attack with 99% confidence, he's fine if he shuts down the business. But if a firewall rule is triggered, that's not a good enough indicator, he said.

Malcolm Harkins, chief security and trust officer at Cymatic, recalled an extreme example of one such sub-par indicator causing a massive business shutdown while in a previous role in another organization.

“We have to think about [AI] from the context of risk.”

Malcom Harkins, Chief Security and Trust Officer at Cymatic

“I had 26,562 systems get offline for a day because of a bad DAT from McAfee 10 years ago,” Harkins recalled. “We have to look at it and go, ‘is the proportion of those types of unintentional errors hopefully (getting) better in AI?’”

Those improvements will happen, just like they have in other systems, Harkins said. But, is the AI so powerful that the impact of the error is bigger because it happens faster?

“We have to think about that from the context of risk,” Harkins said. “And then weigh that against ‘are we getting yield off it?’ Because all automation is supposed to provide is efficiency and effectiveness.”

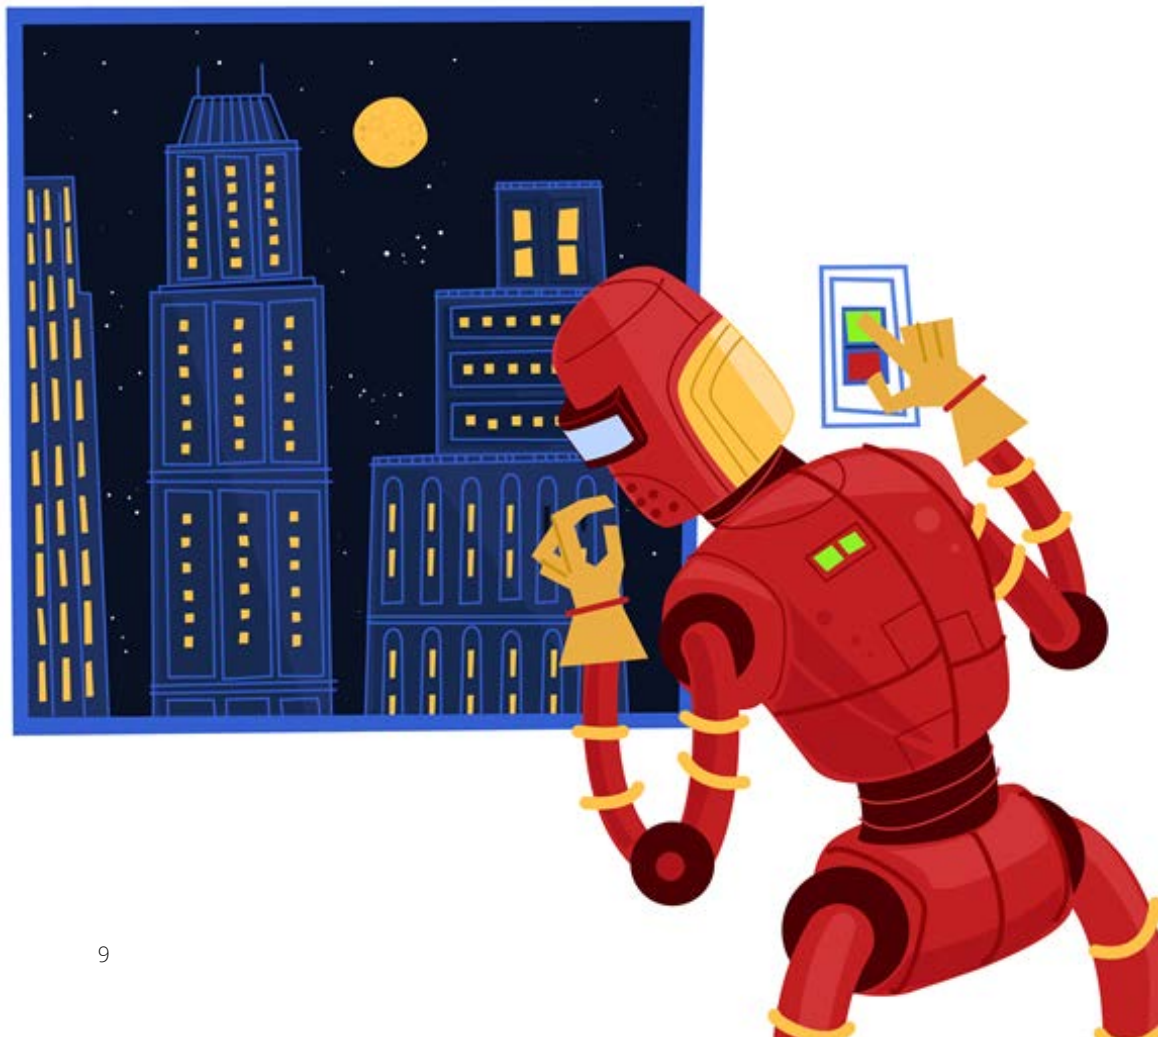

04/05
Find the ROI of AI

According to a 2019 Capgemini Research Institute report*, two-thirds of the 850 executives surveyed said that AI-based tools lowered costs in detecting and responding to security breaches, and three-fourths said AI enabled faster responses, with the time taken to detect threats being shortened by up to 12% on average.

The crux is how these organizations use AI effectively, and if they feel they can depend on the *provability* of the AI tool – meaning, "why did it make the decision it made?" This same Capgemini report demonstrates a larger percentage of implementation among AI elements that serve as indication engines for Curry's "carbon-based" intelligence. On provability, Curry thinks we're a little too sensitive.

"We look at a (self-driving) car, and say, 'We can't trust it, I don't know why it made that decision,'" Curry said. "I've never seen an accident where you could trust any of the drivers for why they made their decisions. But it can be better with AI."

Take the car example. If there's 30-40,000 deaths per year** in automobile accidents, how much lower does it have to be before you don't care? Is it 300? Is it three?

"At what point do you say, 'it's good enough?'" Curry asked.

Along with provability, CISOs also need to consider what time and resources they devote to training machine learning and other vendor AI models for which they're already paying a premium.

"Unless I built it myself, it wouldn't be in my interest to train someone else's AI. I'd rather spend my time finding a vendor whose AI I didn't have to train," Taylor Lehmann, a four-time CISO, said.

You have to decide, Lehmann said, "Is the value of my supervision of that model going to give me a higher efficacy of control? I haven't run into a situation yet where the cost benefit for me as a CISO is so high that I would dedicate a person or a substantial amount of resources to go ahead and train that solution internally."

The vital piece, when entering vendor negotiations, is "can I test it, can I prove that it works, can I show, based on the amount of effort that I put into this, do I get the ROI," Lehmann said.

*Reinventing Cybersecurity with Artificial Intelligence," Capgemini
**"Motor Vehicle Deaths Estimated to Have Dropped 2% in 2019," National Safety Council

“Unless I built it myself, it wouldn’t be in my interest to train someone else’s AI. I’d rather spend my time finding a vendor whose AI I didn’t have to train.”

Taylor Lehmann, Board Member at Health-ISAC

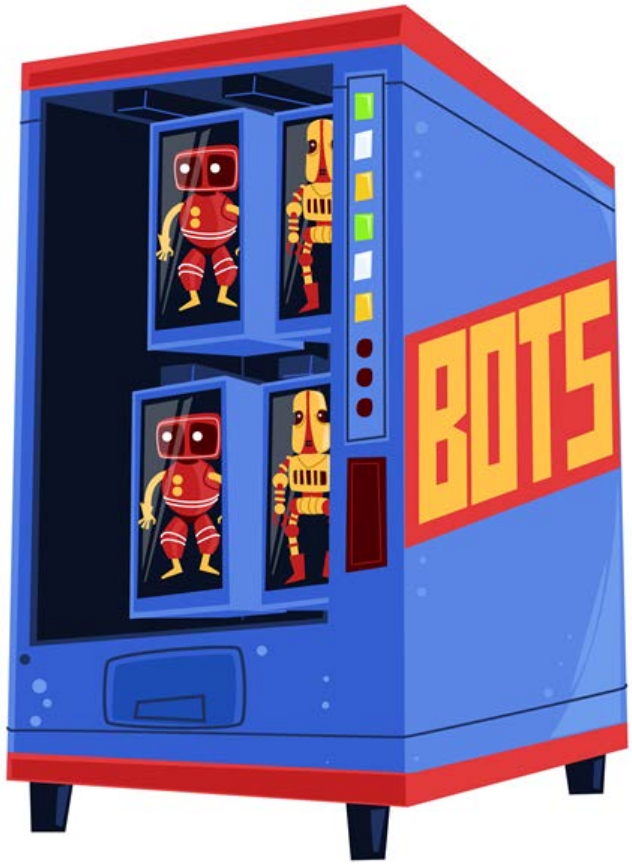

05/05

Think About the Rights: Do You Have Rights to Your Data?

Beyond ROI and provability in security use-cases, how does AI impact your life? Your privacy? How do we define privacy?

“Most privacy regulations are adopting increasingly expansive definitions to account for future forms of collectible data,” Curry said.

“With privacy, the notion of harm is really important. Do you have rights to that data like you would have to your property? If somebody measures it without telling you and puts it to use, have you been hurt?”

Curry defines security to “neophytes” as the “cost to break something.”

“I think of privacy as the cost to obtain information on someone,” he said. “Companies holding data should have a principle that they don’t make it easier for others, and at lower cost, to get information about you — they ‘do no harm’ in that respect.”

Curry suspects that there is a privacy revolution coming, and companies that lean into it will benefit. But somewhere in the future, we’re going to define privacy much more broadly, and we better be getting ahead of it.

With AI and machine learning, there are biases in input, there’s biases in processing, there’s biases in output, and we need to be careful not to inadvertently give people advantages based on those biases.

“Prepare for what exists in legislation today,” Curry said, “but I think we need a bigger perspective long-term.”

“Prepare for what exists in legislation today, but I think we need a bigger perspective long-term.”

Sam Curry, CSO at Cybereason

About Cyber Resilience Think Tank

The Cyber Resilience Think Tank is an independent group of industry influencers dedicated to understanding the cyber resilience challenges facing organizations across the globe, and together, providing guidance on possible solutions.

They define cyber resilience as: “an organization’s capacity to adapt and respond to adverse cyber events — whether the events are internal or external, malicious or unintentional in ways that maintain the confidentiality, integrity and availability of whatever data and service are important to the organization.”

Malcolm Harkins
Chief Security &
Trust Officer
Cymatic

Juan Harmse
Head of Resilience Strategy &
Engagement
ABSA Group

Taylor Lehmann
Board Member
Health-ISAC

Gary Hayslip
CISO
SoftBank Investment Adviors

Sue Lapierre
VP, Information Security
Officer
Prologis

Peter Tran
Head of Global Enterprise
Cyber & Product
Security Solutions
InferSight

Dr. Sam Small
CSO
ZeroFOX

Maurice Stebila
CISO, Digital Cyber Security,
Compliance & Privacy Officer
HARMAN International

Jakub (Kuba) Sendor
Software Engineer
Yelp

James Lugabihl
Senior Director, Global
Security
ADP

Ari Schwartz
Managing Director of
Cybersecurity Services
Venable

Stephen Ward
CISO
Home Depot

Dawie Wemtzel
Head of Forensic
Investigations
ABSA Group

Chris Wysopal
Chief Technology Officer
Veracode

Sam Curry
CSO
Cybereason

Greig Arnold
CISO
Vista Consulting Group

Bill Brown
CSO
ClickSoftware

Marc French
CISO, Managing Director
Product Security Group

Josh Douglas
VP Product Management
Threat Intelligence
Mimecast

Scott Eigenhuis
Associate Director,
Information Security
Illumina

Shawn Valle
CSO, VP
Rapid7

Michael Madon
CEO
Stealth Startup

Claus Tepper
Head of Cyber Security
Operations, ABSA Group

Brian Miller
CISO
Healthfirst

Visit the Mimecast resource center to learn more.

For more insights from the Cyber Resilience Think Tank, visit mimecast.com/ThinkTank