



CYBER DEFENSE

MAGAZINE

eMAGAZINE

SEPTEMBER 2025

In This Edition

A Sea Change is Coming in the Distribution of Cyber Risks: SMBs (Small and Midsize Businesses) in Critical Infrastructure are on the Front Line

Quantum Threat Is Real: Act Now with Post Quantum Cryptography

The CISO of Tomorrow: A Human-Centric Approach to Leading Through AI, Autonomy, and Quantum Uncertainty

Cyber Insurance: 7 Hard Lessons You Need to Prepare For

...and much more...

MORE INSIDE!

CONTENTS

Welcome to CDM's September 2025 Issue-----	10
<i>A Sea Change is Coming in the Distribution of Cyber Risks: SMBs (Small and Midsize Businesses) in Critical Infrastructure are on the Front Line-----</i>	11
<i>Quantum Threat Is Real: Act Now with Post Quantum Cryptography -----</i>	36
By Ann-Anica Christian, SSL2BUY	
<i>The CISO of Tomorrow: A Human-Centric Approach to Leading Through AI, Autonomy, and Quantum Uncertainty-----</i>	43
By Ferris Adi, MBA, CISSP, CISA, CISM, CRISC Cybersecurity Leader & Author, Futurist. CISO Trans Americas Fiber	
<i>Cyber Insurance: 7 Hard Lessons You Need to Prepare For-----</i>	48
By Paul Barbosa, General Manager, Global Cloud Security Business Unit, Check Point	
<i>Adapting Incident Response to Autonomous Agents: Evolving Practices for the Future of Cyber Defense-----</i>	52
By Tannu Jiwnani, Principal Security Engineer, Microsoft	
<i>Adversarial GRC Weaponizing Compliance to Defend against Synthetic Threats -----</i>	57
By Victor D Patterson Sr, AI Cybersecurity Strategist & Founder of DeepSecure™ — Architect of AI Security Doctrine	
<i>The Future of Mobile App Security: Emerging Technologies and Trends-----</i>	62
By Jason Cortlund, Mobile App Security Evangelist at Guardsquare	
<i>More Than Just a Harmless Box: Making the Case for Printer Platform Security -----</i>	66
By Steve Inch, Global Senior Print Security Strategist & Product Management Lead at HP Inc.	
<i>Preparing for Updated HIPAA Security Requirements -----</i>	69
By Erik Eisen, CEO of CTI Technical Services	
<i>Email Threats Are Evolving: Why Low-Tech, AI-Powered Attacks Are Outpacing Traditional Defenses -----</i>	73
By Usman Choudhary, Chief Product and Technology Officer, VIPRE Security Group	
<i>Zero Trust: What Cybersecurity Experts Can Learn from Prisons -----</i>	77
By Nick Kathmann, CISO, LogicGate	

5 Steps to Move Beyond Vulnerability Discovery to Risk Remediation -----	80
By Chaz Spahn, Director of Product Management at Adaptiva	
Agentic AI is Transforming Security – What Enterprise Leaders Need to Know -----	84
By Michiel Prins, Co-Founder & Senior Director, Product Management, HackerOne	
AI is the New Attack Surface -----	89
By Patrick M. Hayes, Author & Security Strategist, Integrated Assurance LLC	
Artificial Intelligence Has Changed What Cyber Defense Means for the Retail Industry -----	94
By Jason Lewkowicz, Managing Director – Cybersecurity, America’s Retail Sector Lead, Ernst & Young LLP	
Branding Beyond the Breach: How Cybersecurity Companies Can Lead with Trust, Not Fear -----	98
By Lucia Barbato, CEO, Ilex Content Strategies	
Building Responsible AI: Staying Ahead in a Regulated World -----	102
By Paul Davis, Field CISO, JFrog	
Wiper Malware: What Federal, State and Local Agencies Must Know to Protect Mission-critical Systems -----	106
By Dirk Schrader, VP of Security Research and Field CISO EMEA, Netwrix	
Cloud-Native, AI-Driven, and Always-On: The Future of Firewall Security -----	110
By Paul Barbosa, Global Cloud Security Business Unit Leader, Check Point Software Technologies	
Man, Machine and Malware -----	114
By Marlene Francis, Research Analyst, IndustryARC	
Threat Intelligence in the Heat of Cyber Warfare -----	121
By Andrew Martin, CEO at DynaRisk	
Deepfakes, Synthetic Media, and Digital Trust: The Cybersecurity Implications of Deepfake Technology and Methods for Detection and Mitigation -----	124
By Joe Guerra, M.S., Software Engineering, CASP+,CCSP, FedITC,LLC	
How Red Teams are Reinventing Cybersecurity for the Age of AI -----	130
By Dr. Darren Pulsipher, Chief Solution Architect, Intel	
Inside the Mind of a Threat Actor: What CISOs Must Learn Before the Next Breach -----	133
By Ahmed Awad (aka nullc0d3), Senior Cyber Threat Intelligence Analyst, Author & Educator	

<i>AI Adoption in Application Security</i>-----	136
By Kaushik Majumder, Director, Sony India Software Centre	
<i>Dormant Access and the Hidden Risk inside your IAM Program</i>-----	141
By Durgaprasad Balakrishnan, Independent Cybersecurity Researcher and Director of Cybersecurity – Identity and Access Management at a Leading Global Fintech Company	
<i>Personal Cybersecurity: Benefits for Individuals and Your Business</i>-----	146
By Paul Pioselli, Founder & CEO, Solace – Truly Personal Cybersecurity	
<i>Shadow APIs: The Silent Backdoor Undermining Application Security</i>-----	151
By Sandeep Dommari, Principal Architect, Ping Identity	
<i>The Rise of AI-Driven Credential Stuffing: Why Identity and Access Management (IAM) Alone Can't Save You</i>-----	157
By Sandeep Dommari, Principal Architect, Ping Identity	
<i>The Future of Privileged Access is Vault-Free</i>-----	162
By Yaron Kassner, Co-Founder & CTO, Silverfort	
<i>How In-Memory Vulnerability Scanning Boosts Enterprise Linux Security</i>-----	167
By Eric Hendricks, Product Manager for the Radar Vulnerability Scanner at TuxCare	
<i>The Vibe Coding Conundrum: Development vs. Security</i>-----	170
By Amir Kazemi, Director, Product Marketing, Cycode	
<i>Creating a Full-spectrum Cyber Workforce for CJADC2</i>-----	173
By Magdalena LoGrande, Cybersecurity Engineering Fellow, Sigma Defense	
<i>Cyber Risk Management for FinTechs in the 21st Century</i>-----	176
By Oreoluwa Joda, Cyber Risk Manager, Fiserv	
<i>Cybersecurity Compliance – Changing the Paradigm with AI</i>-----	184
By David Ramirez, CISO, Broadridge	
<i>Cybersecurity In Critical Infrastructure: Protecting Power Grids and Smart Grids</i>-----	188
By Kehinde Ayano, Assistant Professor of Computer and information Science, Indiana Wesleyan University Marion Indiana USA	
<i>Data Breaches Are the New Normal – Complacency Is the Real Crisis</i>-----	193
By Rob Vann, Chief Solutions Officer at Cyberfort	

<i>Deepfakes and Disinformation: Protecting Truth in the Age of AI-----</i>	198
By Omkar Bhalekar, Sr Network Engineer, Tesla	
<i>Securing the Mission: Why External Cyber Defense is Essential for Government and Education -----</i>	203
By Amit Weigman, Office of the CTO, Check Point Software Technologies	
<i>Erase, Don't Just Delete: How Secure File Wiping Can Stop Insider Threats -----</i>	207
By Neha Sawhney, Marketing Specialist, Stellar Information Technology Pvt. Ltd.	
<i>From Shadows to Spotlight—Why Low-Priority Assets Now Define the Frontline of Cybercrime -----</i>	211
By Pankit Desai, CEO of Sequaretek	
<i>Grappling with a Post-CVE World -----</i>	216
By Tod Beardsley, VP Security Research, runZero	
<i>Guardians of the Factory: Defending Cyber-Physical Systems in Smart Manufacturing-----</i>	219
By Omkar Bhalekar, Sr Network Engineer, Tesla	
<i>How Digital Twins Enhance Grid Security -----</i>	223
By Zac Amos, Features Editor, ReHack	
<i>Breaking Traditional Encryption Protocols: Quantum Computing and the Future of Secure Communications -----</i>	227
By Joe Guerra,M.Ed. CASP+,CCSP, FedITC,LLC	
<i>IGA Is a Journey, not a Destination – Plan Accordingly -----</i>	234
By Niels Fenger, Advisory Practice Director, Omada	
<i>The Perimeter Is Not Enough-----</i>	238
By Parker Pearson, Chief Strategy Officer, Donoma Software	
<i>Passwordless, but Not Riskless -----</i>	243
By Sudhakar Tiwari, Principal Solutions Architect, Zurich	
<i>Privileged Access in the Age of Ransomware--as--a--Service (RaaS) -----</i>	249
By Sandeep Dommari, Principal Architect, Ping Identity	
<i>Reachability and Exploitability-----</i>	253
By Julia Lorenz, Solutions Manager at Xygeni	

<i>Responsible AI in the Cloud: Building A Framework for Trust</i> -----	259
By Marina Bregkou, Principal Research Analyst and Associate Vice President, Cloud Security Alliance	
<i>SASE Architecture Deployment Across the Globe</i> -----	263
By Vishal Gudhka, Senior Network Architect, Versa Networks	
<i>Security Chaos Engineering for CISOs - The Strategic Edge Against Modern Threats</i> -----	268
By Andres Andreu, CISO and COO, Constella Intelligence	
<i>Generative AI Security: Protecting AI Workloads with AWS Cloud</i> -----	274
By Jatinder Singh, Senior Technical Account Manager, Amazon Web Services	
<i>Part 1: Maximizing the Value of LLMs Without Compromising Security: The Identity Risk Behind AI Agents</i> -----	282
By Amit Zimerman, Co-Founder and CPO, Oasis Security	
<i>The Ghost in the Firewall</i> -----	285
By Bekir Tolga TUTUNCUOGLU, CEO & Researcher, TTNC Teknoloji	
<i>The New Face of Cyber Threats: From Scattered Spider to North Korea's Phantom Workforce</i> -----	290
By Richard K. LaTulip – A Field Chief Information Officer at Recorded Future	
<i>The Real Cost of Exposure Remediation: Helping Developers Avoid Burnout</i> -----	295
By Ravid Circus, Chief Product Officer, Seemplicity	
<i>The Rise of AI-Driven Credential Stuffing: Why IAM Alone Can't Save You</i> -----	298
By Sandeep Dommari, Principal Architect, Ping Identity	
<i>The Role of Agentic AI in Proactive Cyber Threat Hunting</i> -----	303
By Nivedita Kumari, Data & AI Customer Engineer, Google	
<i>The Security-speed Myth That's Sabotaging Your Modernization</i> -----	306
By Ayo Akinsanya, CISSP, PMP, ITIL (Cybersecurity Expert)	
<i>The Top Three Emerging AI Threat Tactics Changing the Face of Identity Fraud</i> -----	311
By Ashwin Sugavanam, VP of AI & Identity Analytics	

The AI Arms Race Is Here: How Enterprises Must Weaponize Data and Autonomous Defense to Survive the Next Generation of Cyberattacks ----- 315

By Nic Adams, Co-Founder & CEO, Orcus

Threat Actor Profiling ----- 319

By Diyar Saadi, Computer Security Researcher, Independent Researcher

Modern Risk, Modern Response: Federal Cybersecurity Needs a Compliance Wake-Up Call ----- 325

By Peter O'Donoghue, CTO at Tyto Athene, and Gaurav Pal, CEO at stackArmor (a Tyto Athene company)

Video Monitoring Security ----- 330

By Milica D. Djekic

When Theory Meets Reality: The Fatal Flaws in Traditional Incident Response ----- 333

By Andy Lunsford, CEO and Co-Founder, BreachRx

Where Do Leadership and Technical Teams Disagree on AI-Driven Network Security? ----- 337

By Mitch Densley, Principal Solutions Architect at Opengear

@MILIEFSKY

From the

Publisher...



On the heels of our very successful participation in Black Hat, we have published numerous Innovator Spotlight interviews and have connected with many new participants in our Cyber Defense Media Group network. It has been a very fulfilling and gratifying time of expansion and dedication to our mission.

In addition to this September issue of Cyber Defense Magazine covering an expanding spectrum of cybersecurity issues, we recognize that the velocity of developments continues to increase. With particular emphasis on uses and abuses of artificial intelligence, these developments occur more and more frequently. That brings home the importance of timely reporting and analysis through our eMagazine and online coverage.

If you're curious about where AI is headed — and what it means for our future — check out my new book, *The AI Singularity: When Machines Dream of Dominion*, available now on Amazon: <https://amzn.to/4dPyakN>

We are also preparing for the Cyber Defense Conference coming in October. You can see more detail at <https://cyberdefenseconferences.com>

Our flagship top awards programs for 2025 are still open for nominations:

- 🏆 Top Global CISOs
- 🏆 Top InfoSec Innovators
- 🦄 Black Unicorn Awards

These are global platforms that recognize true cybersecurity leadership and innovation. Whether you're a Chief Information Security Officer, a disruptive cybersecurity startup, or a fast-scaling company poised to become the next unicorn, this is your chance to shine on the world stage.

Nominate yourself, your company, or someone you admire at: 🗑️ <https://www.cyberdefenseawards.com>

Winners will be recognized at **Cyber Defense Conference 2025**, taking place October 28-29 in Orlando, Florida — an exclusive, high-trust gathering of 300+ CISO award winners and top cybersecurity leaders. Learn more at <https://cyberdefenseconferences.com>.

Stay sharp, stay secure — and remember:

Cybercriminals never sleep. Neither can your cyber defense.

Warmest regards,

Gary S. Miliefsky

Gary S. Miliefsky, fmDHS, CISSP®
CEO/Publisher/Radio/TV Host

P.S. When you share a story or an article or information about CDM, please use #CDM and @CyberDefenseMag and @Miliefsky – it helps spread the word about our free resources even more quickly



@CYBERDEFENSEMAG

CYBER DEFENSE eMAGAZINE

Published monthly by the team at Cyber Defense Media Group and distributed electronically via opt-in Email, HTML, PDF and Online Flipbook formats.

EDITOR-IN-CHIEF

Yan Ross, JD

yan.ross@cyberdefensemagazine.com

ADVERTISING

Marketing Team

marketing@cyberdefensemagazine.com

CONTACT US:

Cyber Defense Magazine

Toll Free: 1-833-844-9468

International: +1-603-280-4451

<https://www.cyberdefensemagazine.com>

Copyright © 2025, Cyber Defense Magazine, a division of
CYBER DEFENSE MEDIA GROUP

1717 Pennsylvania Avenue NW, Suite 1025

Washington, D.C. 20006 USA

EIN: 454-18-8465, DUNS# 078358935.

All rights reserved worldwide.

PUBLISHER

Gary S. Miliefsky, CISSP®

Learn more about our founder & publisher at:

<https://www.cyberdefensemagazine.com/about-our-founder/>



13 YEARS OF EXCELLENCE!

Providing free information, best practices, tips, and techniques on cybersecurity since 2012, Cyber Defense Magazine is your go-to-source for Information Security. We're a proud division of Cyber Defense Media Group

[CYBERDEFENSEMEDIAGROUP.COM](https://www.cyberdefensemediagroup.com)

[MAGAZINE](#)

[TV](#)

[RADIO](#)

[AWARDS](#)

[PROFESSIONALS](#)

[WIRE](#)

[WEBINARS](#)

[CYBERDEFENSECONFERENCES](#)

Welcome to CDM's September 2025 Issue

From the Editor-in-Chief

We are delighted to publish this September 2025 issue of Cyber Defense Magazine. In addition to the growth of our readership and breadth of topics, we are enjoying a great deal of positive feedback from cybersecurity professionals and the organizations which depend on their services to maintain cyber resilience and sustainability.

Among the participants in the 16 sectors of critical infrastructure, there is a heightened appreciation for the necessity of assuring that the supply chains of both products and services are dependable. This awareness is reflected in growing requirements for small and midsize businesses (SMBs) to implement effective cybersecurity practices and also to secure insurance coverage on a greater scale.

This month, for the first time, we are publishing an editorial to highlight these trends and the actions required of SMBs in critical infrastructure to reapportion cyber risks. We contemplate an initial series of three columns, starting with banking and the relationships of SMBs to their financial services. Second will be health care and health care finance, where many SMBs are designated as "business associates" of HIPAA-regulated organizations. Third will be defense and aerospace industries, where SMBs provide many critical products and services.

In shedding light on these trends, and in coordination with our outreach to academic institutions offering cybersecurity courses, we intend to play an active role in bringing together those with cybersecurity needs with the sources of fulfilling them. As an adjunct resource for education and training, Cyber Defense Magazine is an important contributor to achieving a balance of supply and demand.

A reminder is in order that CDM is a non-political publication; we do not endorse or oppose any political views expressed by our authors.

Wishing you all success in your cybersecurity endeavors,



Yan Ross
Editor-in-Chief
Cyber Defense Magazine

About the US Editor-in-Chief

Yan Ross, J.D., is a Cybersecurity Journalist & U.S. Editor-in-Chief of Cyber Defense Magazine. He is an accredited author and educator and has provided editorial services for award-winning best-selling books on a variety of topics. He also serves as ICFE's Director of Special Projects, and the author of the Certified Identity Theft Risk Management Specialist ® XV CITRMS® course. As an accredited educator for over 20 years, Yan addresses risk management in the areas of identity theft, privacy, and cyber security for consumers and organizations holding sensitive personal information.

You can reach him by e-mail at yan.ross@cyberdefensemagazine.com





A Sea Change is Coming in the Distribution of Cyber Risks: SMBs (Small and Midsize Businesses) in Critical Infrastructure are on the Front Line

Part 1: The Banking Case

In this first of a new series of editorials from Cyber Defense Magazine, we focus on banking relationships of small and midsize businesses. There are two aspects: as suppliers to financial institutions and as customers using working capital loans and other banking services.

The context is the general shift of risks away from larger customers of SMBs and other organizations with which SMBs have business relationships. Future editorials will focus on health care and health care finance, and defense industries.

Historically, SMBs have treated cybersecurity as a cost and have been reluctant to commit funds to the various types of cyber risk management. There are several reasons for this stance. SMBs tend to consider themselves too small for cyber criminals to attack. They ignore the fact that ransomware and data breaches, among others, are easier for criminals to access and turn against businesses. They don't consider the existential risk of being unable to pay to restore services and resume normal operations.

The numbers are quite different, however. In the past couple of years, over 10 million SMBs have suffered cyber attacks; about a half-million each year just go out of business entirely as a result. The risk of such failures does not fall on the affected SMBs alone.

The rub comes where their customers and clients become unwilling to shoulder this risk of non-performance. It has become more and more common for the customers, especially in the supply chains of critical infrastructure, demand to have their SMB suppliers demonstrate that they have implemented cyber risk management measures. In addition, they are increasingly requiring proof of cyber risk insurance. Should the SMB supplier be unable to fulfill the terms of a contract, the purchaser needs assurance that there is a deeper pocket available to help pay for delays and replacements.

Once this set of dynamics is recognized, it's easy to see how the result is the shifting of risk from the buyers and insurers to the SMBs. What has been an optional expense is moving in the direction of becoming a requirement to stay in business.

This first column focuses on the relationships of SMBs with financial institutions. Of course, there are SMBs providing products and services to banks and other financial institutions, including insurance companies and securities firms. There is no question that the timeliness and accuracy of SMB suppliers are paramount in fulfilling such contracts.

The broader and even more vulnerable aspect of their relationships tend to be on the customer side, where the SMBs typically depend on banking services and working capital loans to stay in business. Any cyber event which impairs the ability of the SMB-customer to comply with the terms of such a loan becomes a problem for the bank.

It's not just the banks which must take a fresh look at the cyber vulnerabilities of their credit customers. They are regulated heavily by both State and federal agencies. This regulation is carried out through rigorous examinations. Weak borrowers discovered in regulatory examinations can result in both write-offs and additions to loan loss reserves and even civil money penalties.

Notably, in a recent FDIC report on examinations, only 2 pages out of 80 pages are devoted to cybersecurity. Similar factors are in play at the Small Business Administration, which is responsible for guaranteeing SBA loans. However, there can be little doubt that as cyber attacks grow, and ransomware becomes a greater threat to SMBs and their ability to service their loan obligations may be impaired, that the regulators will impose more stringent requirements.

But it is certain that stricter standards will be observed as the ease of cyber attacks grows and the vulnerabilities of SMBs continue. The only effective response must be for SMBs to undertake cybersecurity measures. The other choice is to risk losing not only business opportunities but the entire operation as a going concern.

That brings us full circle to the mission of Cyber Defense Magazine and our parent Cyber Defense Media Group: to provide for CISOs and all organizations using their services to have access to actionable intelligence in support of effective cybersecurity programs.

We stand ready to respond to inquiries from our readers in this regard. In addition to our internal resources, we are fortunate to have hundreds of authors and sponsors, and their organizations available as resources to provide valuable advice to all.

Among these resources, there are services available which have offered to perform deeply discounted cyber risk assessments, for those SMBs wishing to get started but don't know where to begin. If this offer appeals to your company, please check in with us at www.cyberdefensemagazine.com.

Next month, we will write more about SMBs in critical infrastructure, with focus on health care providers and sources of payment for health care.

Cyber Defense Magazine is pleased to invite comments and suggestions on this new facility for our growing readership.



SPONSORS



"NightDragon Security is not looking to invest in 'yet another endpoint' solution or falling for the hype of 'yet another a.i. solution', it's creating a unique platform for tomorrow's solutions to come to market faster, to breathe new life into a stale cyber defense economy"

-David DeWalt

Managing Director and Founder NightDragon Security

ADVISE

WE DELIVER SOUND ADVICE AS YOUR FINANCIAL PARTNERS

INVEST

WE ARE FLEXIBLE INVESTORS ACROSS ALL STAGES OF GROWTH TO PRE-IPO

ACCELERATE

WE HELP OUR COMPANIES ACCELERATE THEIR GROWTH THROUGH STRATEGY TUNING AND RELATIONSHIP BUILDING

www.nightdragon.com



AN INTEGRATED GLOBAL, STAGE-AGNOSTIC CYBERSECURITY
INVESTMENT PLATFORM SPANNING SEED THROUGH
GROWTH.

The first dedicated cybersecurity venture firm in the world

About us

For 20 years, AllegisCyber Capital has offered a proven investment platform that actively engages with the most promising entrepreneurs in cyber and delivers the market access, team building, and operating experience essential to their success.

BUILDING AND FUNDING THE MOST INNOVATIVE COMPANIES IN CYBER



www.allegiscyber.com



hello@venturescope.com
www.venturescope.com
@venturescope



VentureScope®

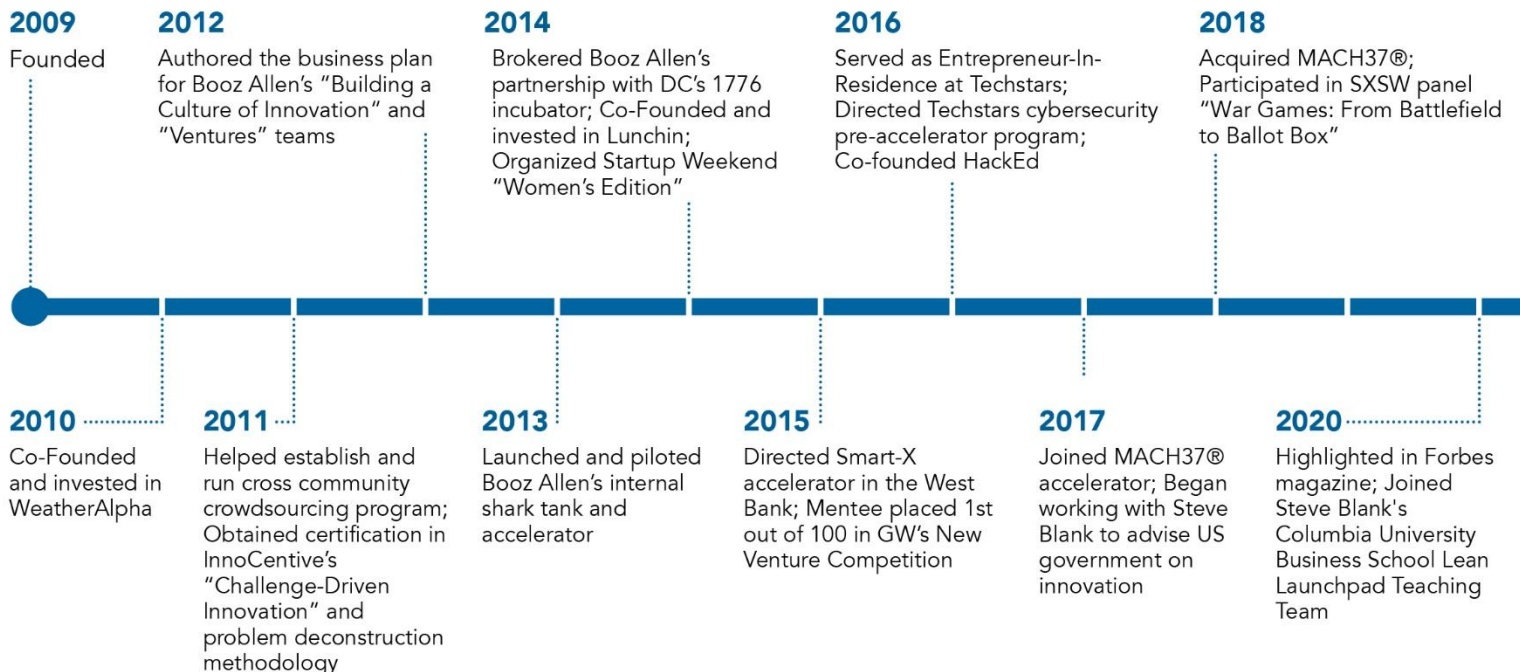
STRATEGY · DEEP TECH · INVESTMENT

VentureScope® works with creative entrepreneurs, venture capital investors, and large private and public sector organizations around the world that are trying to solve interesting problems. Our team specializes in problem deconstruction and framing, product development, business model refinement, go-to-market strategies, build-buy-partner decisions, strategic partnerships, investment and growth analysis, and a variety of innovation methodologies. Whether you're a budding entrepreneur, a scrappy startup, an experienced investor, or an established organization developing a new service or capability, we will not only advise you on what to do, but work as part of your team to apply our recommendations.

Our team has over 60 years of combined experience launching new business ventures, investing in promising startups, running startup accelerators, teaching and providing strategic innovation and general management consulting services to large private and public sector organizations. We own and operate the MACH37 Cyber Accelerator®. We're on the pulse of emerging and over-the-horizon technology, and are tracking their growth and development against important industry problems to inform our dealflow and give you exceptional advice.

Expertise

LEAN STARTUP METHODOLOGY
BUSINESS MODEL STRATEGY
PROBLEM DECONSTRUCTION & FRAMING
PRODUCT DEVELOPMENT
GO-TO-MARKET STRATEGY
REVENUE GENERATION
TECHNOLOGY SCOUTING & INVESTMENT DEALFLOW
BUILD-BUY-PARTNER DECISIONS
INVESTMENT & GROWTH ANALYSIS
STRATEGIC PARTNERSHIPS
CHALLENGE-DRIVEN & OPEN INNOVATION
INNOVATION PIPELINE DESIGN & IMPLEMENTATION
CREATIVITY & STRATEGIC FACILITATION
INSTRUCTIONAL DESIGN & EXPERIENTIAL TRAINING
HUMAN PERFORMANCE





“Built on passion and expertise, Altitude Cyber delivers strategic advisory services specifically tailored for founders, investors, startups, and their boards. Our unique approach fuses strategic insight with financial acumen to help your company soar to new heights.”



Dino Boukouris

Managing Partner, Altitude Cyber

Guiding cybersecurity businesses globally through every stage of growth with tailored advisory services for founders, CEOs, investors, and boards.



Founders & CEOs

Altitude is your trusted advisor throughout your entrepreneurial journey. We guide you as you grow your business, navigate fundraising processes, construct advisory boards, plan your long-term exit strategy, develop strategic relationships with key partners and investors, and more.



Investors

We offer a range of strategic advisory services to support your existing portfolio companies, as well as your potential investments or acquisition targets. Our solutions are tailored to fit your needs, with flexible engagement models that align incentives to maximize outcomes.



Boards

We provide in-depth strategic advisory services, tailored to align with the evolving needs of growing businesses. Our support includes strategic business and corporate development, mergers & acquisitions, corporate finance, long term exit planning, advisor selection, and more.

Firm Highlights

Decades of experience as world class operators and advisors

Highly curated research and thought leadership on strategic activity in the cyber market

Deep industry relationships and partnerships across strategic and financial partners

Cyber Network



15,000+

Cyber
Executives



3,000+

Investors



1,000+

CISOs

Cyber Knowledge

4,500+

Company
Tracker

3,000+

M&A Transactions

8,500+

Financing Transactions

Extensive, global relationships with cyber executives, investors, CISOs, policy influencers, and service providers

Altitude Cyber, LLC | www.altitudecyber.com

For inquiries or further information please contact Altitude Cyber at: dino@altitudecyber.com


We monitor the
DARKWEB
so that your
BUSINESS has
no stops



Jericho Security uses AI to fight AI in a new frontier of cybersecurity

Cyber Threats Evolve—So Should Your Defense

Phishing attacks are no longer generic—they're targeted, adaptive, and constantly evolving. Cybercriminals are leveraging dark web data, deepfake technology, and AI-driven social engineering to bypass traditional defenses.



How it Works Defense That Learns. Security That Wins.

1 Hyper-personalized Phishing Simulations

Mimic and keep pace with real-world spearphishing tactics using **dark-web intelligence**, **social engineering**, and **deepfake deception** to test and prepare employees across **multiple channels** (email, text, audio).

2 Adaptive Security Training Videos


Dynamically customize training based on employee risk profiles and attack patterns, ensuring tailored, effective learning experiences.

3 Automated Threat Remediation

Detect, analyze, and take actions on phishing attempts instantly, feeding data back into the system to strengthen defenses over time.

4 Seamless Security Stack Integration

Works with existing **SIEM**, **email security**, and **compliance** platforms, enhancing interoperability and real-time threat intelligence sharing.



Secure Agentic AI with Cequence

AI agents' autonomous decision-making capabilities present unique security challenges that traditional measures may not fully address. Fortify your AI use against emerging threats with Cequence.



**Go to cequence.ai/assessment
or scan the QR code
to start a free assessment
of your vulnerabilities.**





CYBER DEFENSE — MAGAZINE —

WHERE INFOSEC KNOWLEDGE IS POWER



www.cyberdefensetv.com

www.cyberdefenseradio.com

www.cyberdefenseawards.com

www.cyberdefenseconferences.com

www.cyberdefensemagazine.com




Securing identities at every interaction

Seamless, intelligent, centralized authorization to better secure the modern enterprise in the age of AI

- | Privileged Remote Access
- | Secure Credentials
- | Privilege & Entitlement Elevation
- | Identity Threat Protection
- | Identity Governance

Learn more about how to secure all human and machine identities with Delinea.

We're On It






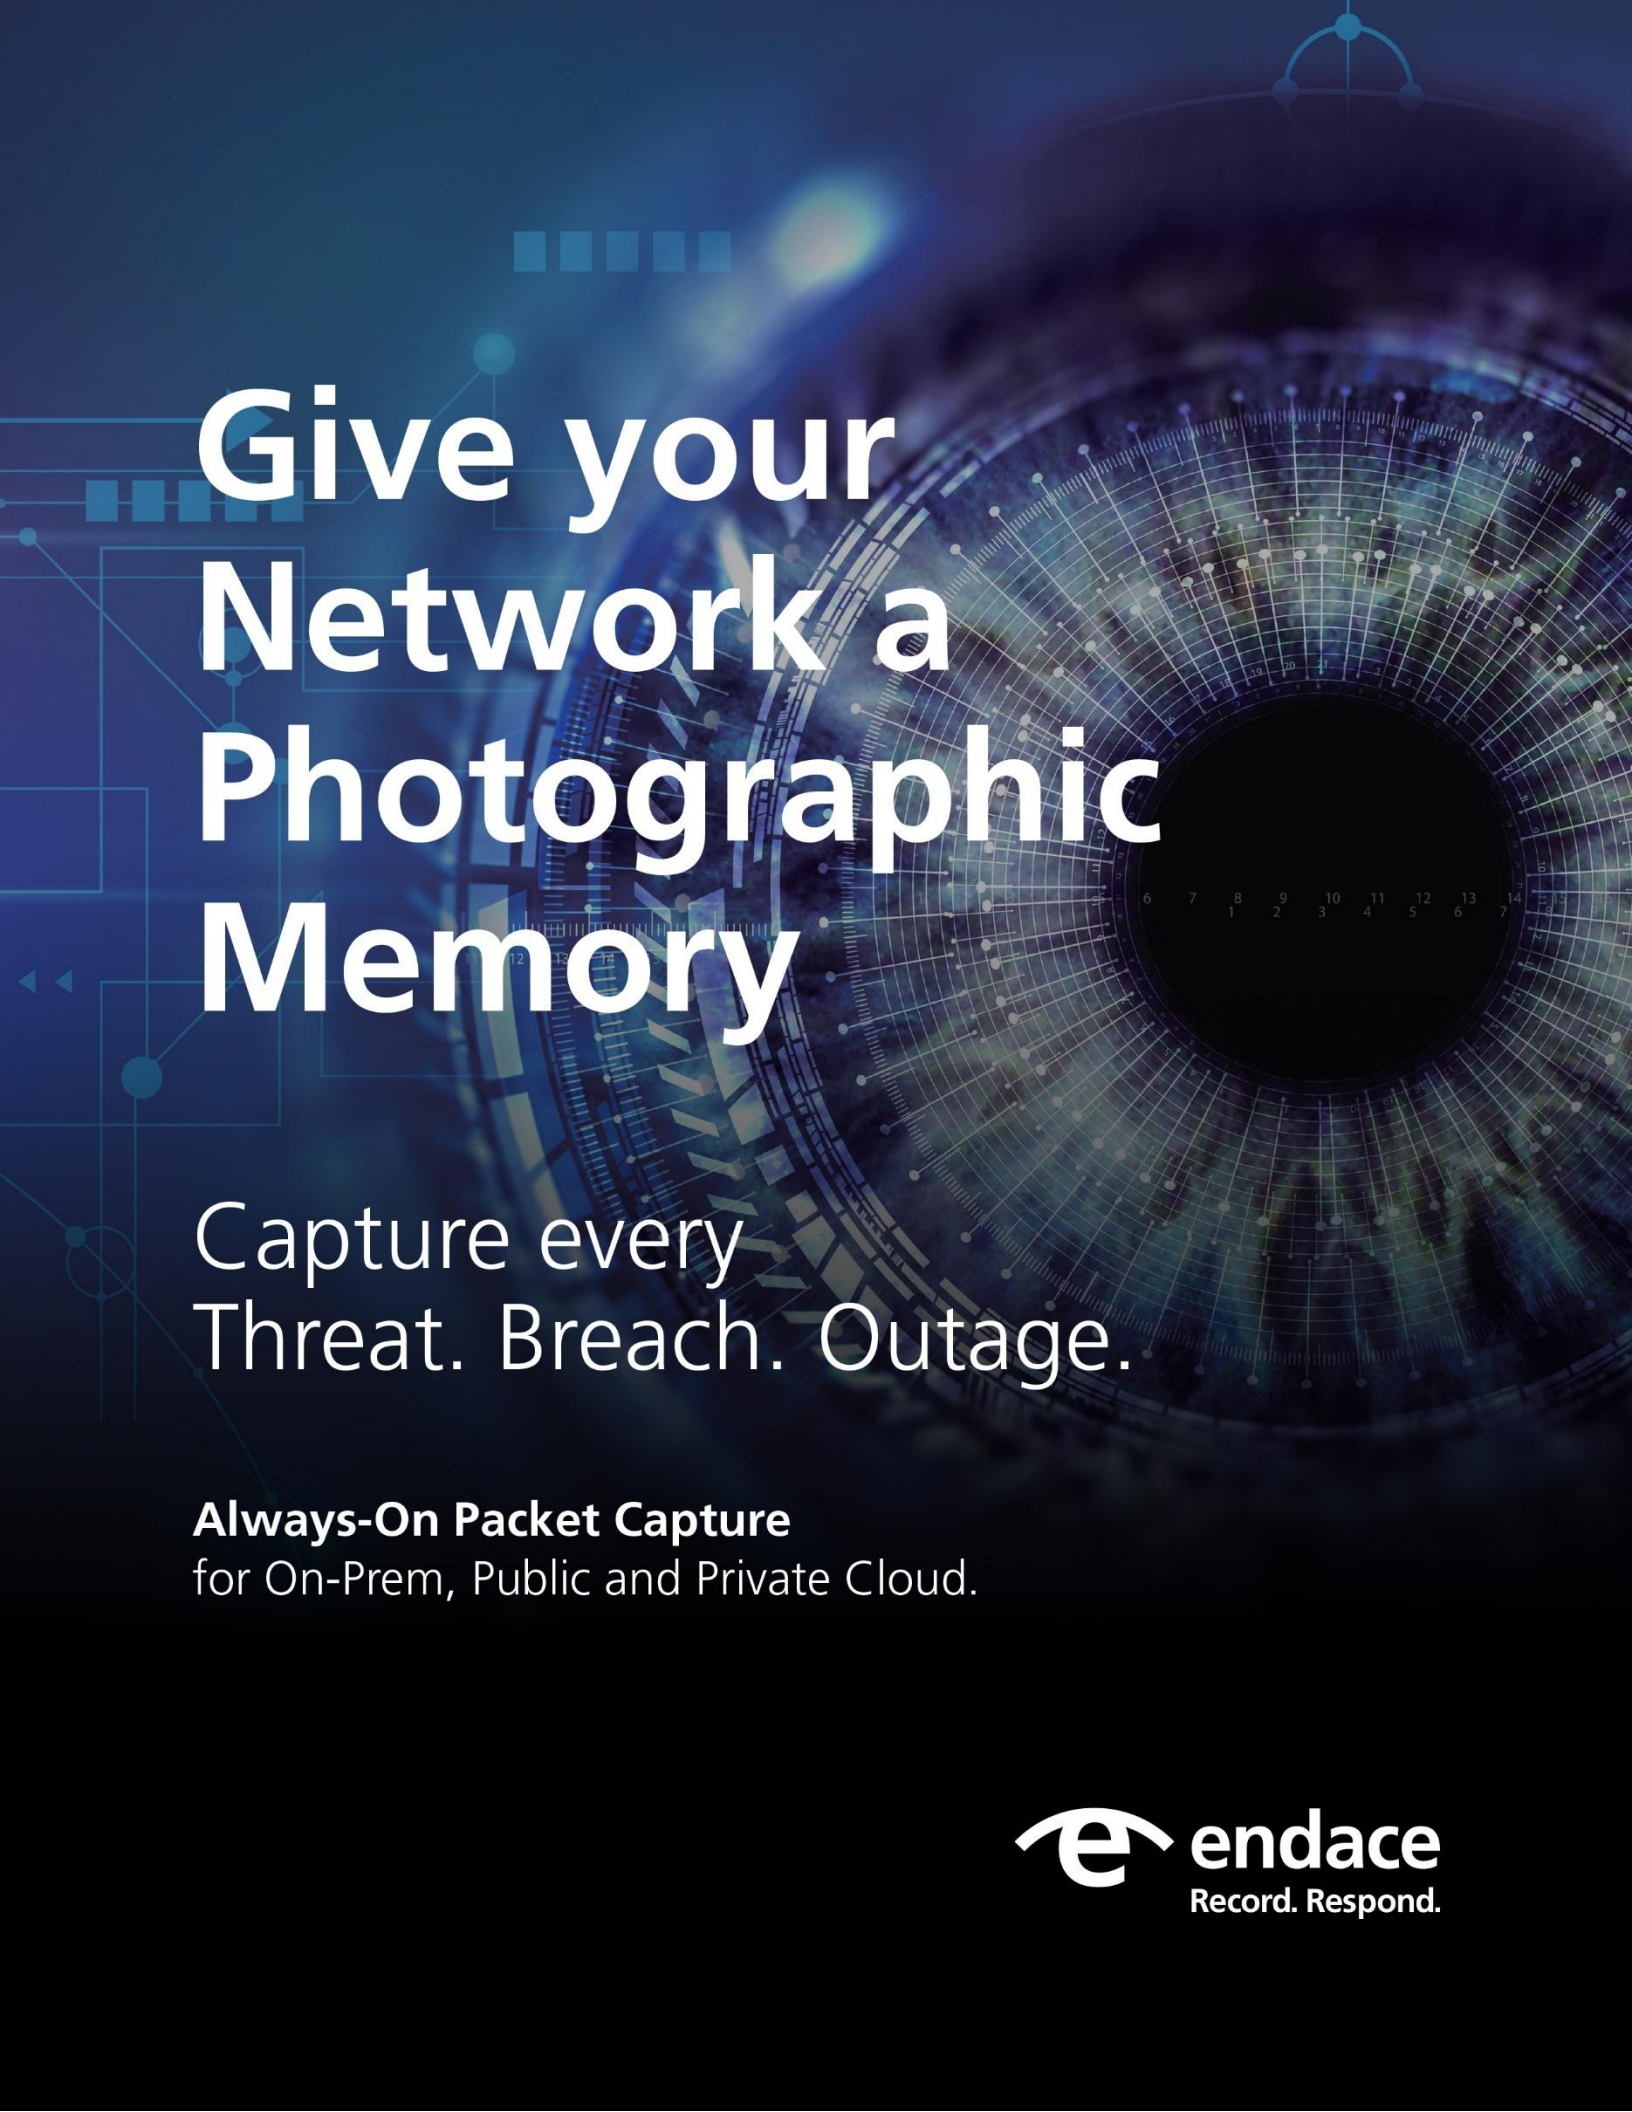
DATA SECURITY FOR THE AI ERA

SAAS | IAAS | FILE STORAGE

- + DSPM
- + Data discovery & classification
- + AI security
- + Data-centric UEBA
- + SSPM
- + Data access governance
- + MDDR

Learn more at www.varonis.com.





Give your Network a Photographic Memory

Capture every
Threat. Breach. Outage.

Always-On Packet Capture
for On-Prem, Public and Private Cloud.

 **endace**
Record. Respond.

JUCY is the Sandbox They Hope You Never Discover

It's time to rethink what's possible!

JUCY Sandbox is the first interactive sandbox to combine dynamic behavioral execution with AI-powered genomic code analysis, running in parallel to catch threats others miss. And unlike traditional sandboxes, JUCY includes hypervisor-based unpacking — invisible to malware and immune to anti-VM evasion.

Developed for the U.S. Intelligence Community and now available for enterprise security teams, JUCY catches new threats designed to evade detection months ahead of other solutions.

Unlike conventional sandboxes that rely primarily on surface-level indicators, JUCY performs deep bytecode analysis to identify malicious code regardless of obfuscation techniques. This groundbreaking approach enables security teams to detect sophisticated malware variants, zero-day exploits, and supply chain malicious insertions that traditional tools fail to recognize.

JUCY works by detonating suspicious files in a secure environment while simultaneously conducting genomic code analysis at multiple levels. The system maps the genetic structure of malicious code, allowing it to identify related malware families even when they've been substantially modified. This function-level detection maintains effectiveness against adversaries who regularly recompile or disguise their code.

Operating at the hypervisor level, JUCY remains invisible to malware, effectively defeating sandbox-aware threats that attempt to evade analysis. The platform's comprehensive memory scanning capabilities also enable it to detect fileless malware and sophisticated memory-resident implants that never write to disk.



www.unknowncyber.com

Application Security, **Reality check.**

Breaking some **myths** about application security



The myth of **simplicity**

Any integration of an open-source library introduces more than 70 additional sub-dependencies.



The myth of **sound analysis**

The application layer is beyond just the code being developed and covered by static scanners, leaving risk valid and unmonitored.



The myth of **accuracy**

Trusting in accuracy without context is a fallacy. More than 90% of alerts are false, generating pure noise.



The myth of **collaboration**

Security tools are never “loved by developers”. Engineering appreciates accuracy, thorough research and professionalism.

The **Application Security Gap** is Growing



Vulnerability backlogs explode as code scales, but developers' capacity to fix stagnates—
Driven by **a lack of clarity and context** to triage and fix issues

Precious engineering time drained by **manual triaging** and **complex remediation steps**

250

Developers

x

\$150K

Average annual
engineer salary

x

5%

Time engineers
spend on security

=


\$1.875M

Added expense
to security

Deloitte.



Ready to build
resiliency? Scan
to get started.



DELOITTE CYBER SERVICES

Do more than defend

In an increasingly open and connected world, cybersecurity is at the core of business success.
Because when you're secure, you can better navigate uncertainty.
When you're prepared, you can turn challenge into opportunity.
And when you're resilient, you can focus less on defending and more on driving forward.

Together, we'll make cybersecurity the core of your success, with the breadth and depth of cyber solutions you need, when you need them.

Ready to build resiliency? Go to Deloitte.com/us/DoMoreThanDefend

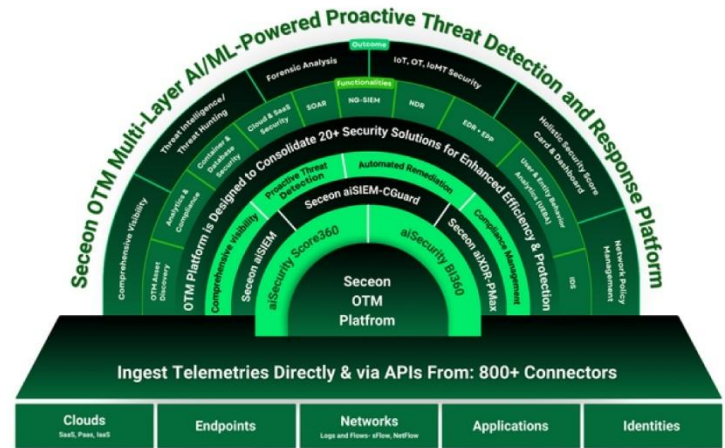
The Current: Trusted news & views
on cybersecurity

Scan to read the latest issue and subscribe >





The Cybersecurity Game Changer



Five Features That Set Us Apart from the Rest

1. Real-Time Insights Through Flows, Logs, & Identities, (Not Just Logs)

Unlike competitors who rely solely on logs, Seceon harnesses the power of network flows, applications logs, & OS logs & identities. Logs offer a limited, after-the-fact snapshot. Flows provide a complete, real-time picture of network activity, empowering you to detect and respond to threats faster and more effectively.

2. Strategic Placement in the Network

Seceon's solution is strategically positioned to sit beside the network, not in the way. This means:

- Comprehensive visibility of north-south (inbound/outbound) and east-west (lateral) traffic.
- The ability to detect hidden "cross-talk" within your network.
- Faster detection and response—because in cybersecurity, every second counts.

3. Smarter Data Collection and Enrichment

Our collectors are designed for efficiency and precision.

- They extract only the relevant data from packet flows.
- Unnecessary information is discarded, leaving you with enriched, actionable intelligence without the noise.
- This streamlined approach ensures better performance and sharper insights than competitors like Huntress and Arctic Wolf.

4. Competitive, Transparent Pricing

Seceon offers enterprise-grade protection at a price point that scales with your business. No hidden fees, no complicated structures—just straightforward pricing that delivers unbeatable ROI.

5. Flexible, Agnostic Connectors

Our platform is built for compatibility:

- Seceon works seamlessly with any environment, no matter the vendor or system.
- Need a new connector? We're known for going above and beyond to integrate with even the most unique setups.

Why Choose Seceon?

Seceon is designed to outpace the competition, offering comprehensive visibility, proactive threat detection, automated remediation & continuous compliance, and flexible integration—all at the lowest TCO saving end-customers more than 40% of the cost for a comparable solution.

Others talk about the platform approach and still come up with multiple kludged together products that lack the common content& situational awareness. Which are easily bypassed by threat actors.

Take Action Now:

Ready to experience the difference?
Let us show you how Seceon redefines cybersecurity:

Visit www.seceon.com
Schedule a Demo Today!



Company HQ:
Westford, MA



Contact Us:
<https://seceon.com/contact-us/>

OFFERING SERVICES

CLIENTS IN OVER
50 COUNTRIES

GROWING

WITH MORE THAN
3 THOUSAND
SECURITY PROFESSIONALS

GLOBAL PRESENCE

OVER
50 THOUSAND
CLIENTS ENROLLED

STRATEGIC ALLIANCE

WITH PR SCIENCE
TECHNOLOGY AND RESEARCH
TRUST & POLYTECHNIC
UNIVERSITY OF PR



PREFERRED PARTNER



SAN JUAN
PANAMA CITY
FT. LAUDERDALE
MEXICO CITY
SAO PAULO
SANTIAGO
BOGOTA
MADRID
MELBOURNE



DATATRIBE

GLOBAL CYBER FOUNDRY THAT INVESTS IN AND CO-BUILDS THE NEXT GENERATION OF CYBERSECURITY AND DATA SCIENCE COMPANIES.

Where the World's Best Come to Build Dominant Companies



About us

DataTribe is a startup foundry that invests in and co-builds world class startups focused on generational leaps in cybersecurity and data science. Founded by leading investors, startup veterans and alumni of the U.S. intelligence community, DataTribe commits capital, in-kind services, access to an unparalleled network, and decades of professional expertise to give their companies an unfair advantage.

BUILDING AND FUNDING THE MOST INNOVATIVE COMPANIES IN CYBER



BLACKCLOAK

EN|VEIL

DRAGONS

CodeDx

refirm labs

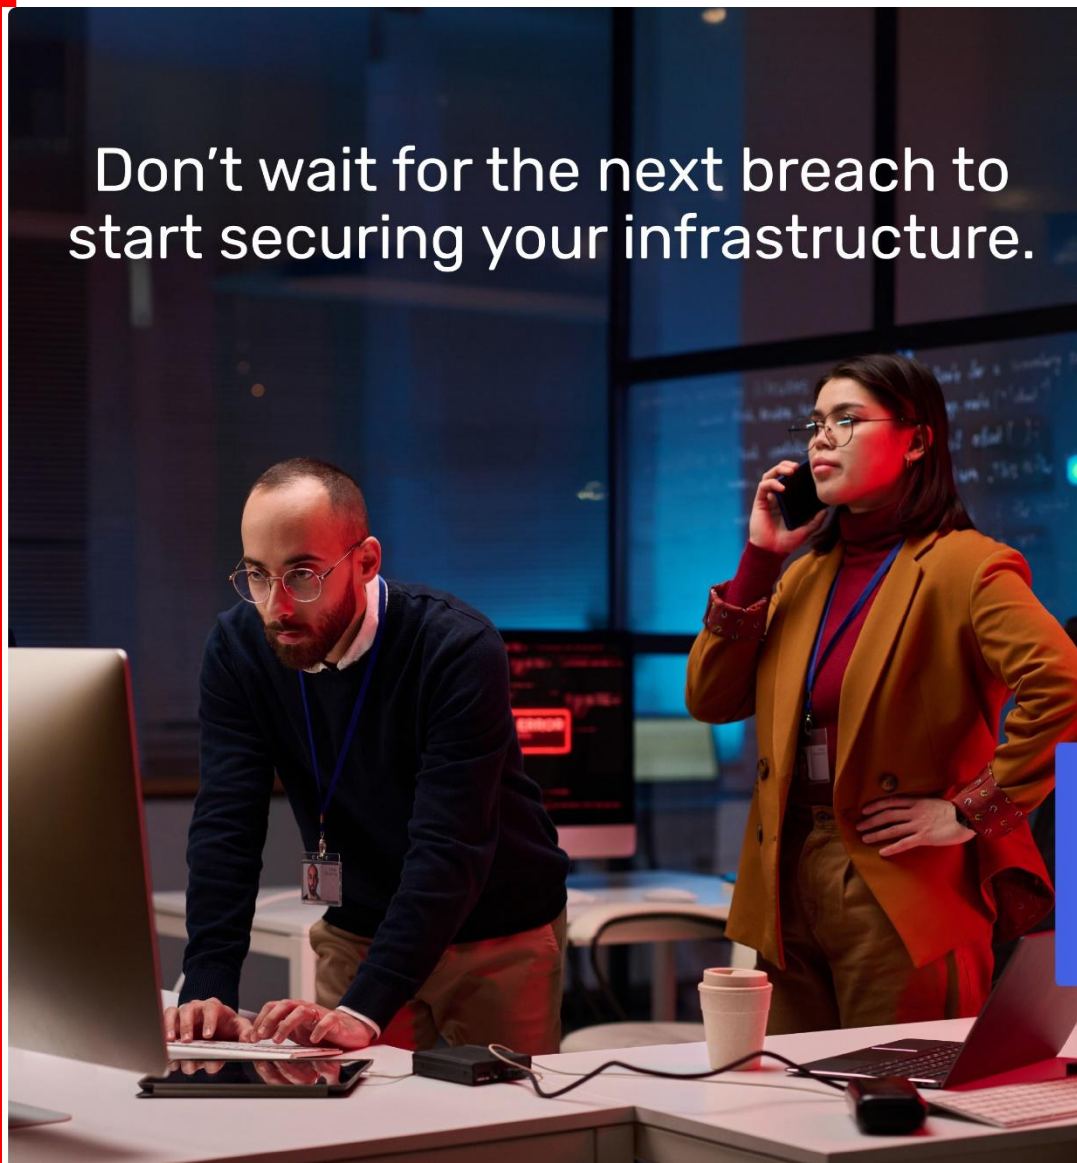
FRENOS

BalanceTheory

fianu

www.datatribe.com





Don't wait for the next breach to start securing your infrastructure.

Strengthening Your Security Post-Breach with Horizon3.ai

In today's rapidly evolving threat landscape, responding after a breach isn't enough. Horizon3.ai leads with a proactive defense approach through its cutting-edge autonomous penetration testing platform, **NodeZero™**.

With continuous cyber risk assessments, NodeZero adapts as your digital environment changes, ensuring long-term protection. NodeZero identifies and prioritizes vulnerabilities, providing actionable remediation strategies. Enhance your security posture, streamline stakeholder communication, and align defenses with organizational goals.




Visit **www.Horizon3.ai** to download our white paper on *Enhancing Cybersecurity Post-Breach* and learn how NodeZero can safeguard your organization against threats.



HORIZON3.ai

THE EXECUTIVE GUIDE TO WINNING AT CYBERSECURITY



Want a Safer, Smarter Cyber Security Strategy?

- Start with the book
- Continue with experts.

Available at
amazon




<https://tinyurl.com/vCISO-Playbook-for-SMBs>



ARTICLES

A hand holding a pen is positioned over a spiral-bound notebook on a wooden desk. To the left of the notebook is a white computer keyboard. In the background, a large book is visible. The entire scene is overlaid with a semi-transparent grey banner containing the word 'ARTICLES' in large, bold, black capital letters. The background image has a soft, blue-tinted aesthetic with some digital network-like patterns.



Quantum Threat Is Real: Act Now with Post Quantum Cryptography

Why Businesses Must Upgrade to Quantum-Safe Encryption Before It's Too Late

By Ann-Anica Christian, SSL2BUY

Cybersecurity has always had to keep pace with the evolution of cyberattacks. These attacks started gaining prominence in the late 80s, in line with the spread of internet access. Initially, viruses, worms and different types of malware were configured to target laptops and desktops. These attacks over time evolved into SQL injection, web-based and social engineering attacks. DDoS attacks followed, and so did Zero Day attacks. Cybersecurity evolved accordingly. Today, the emergence of AI has led to the development of next-gen security platforms. These platforms include Secure Access Service Edge (SASE), XDR and more.

The idea behind cybersecurity is to stay one step ahead of evolving cyberattacks. As the era of quantum computing dawns, it is necessary to be prepared. The [Quantum Computing Cybersecurity Preparedness Act was passed in 2022](#). The focus of this Act is on federal agencies and asking them to prepare against quantum computing attacks. But it makes sense for businesses to be ready as well. [48% of organizations](#) don't have security readiness to address quantum computing threats.

Not a good place to be from a cybersecurity perspective.

Shor's Algorithm

Why are cybersecurity professionals so scared of quantum threats? The reason traces back to the Shor Algorithm. In 1994, Peter Shor, an American theoretical computer scientist, discovered an algorithm. The purpose of this algorithm was to accelerate the factoring of large numbers.

This development was a problem.

Computers that are commonly used today are hard-pressed to factor huge numbers. This is why modern encryption systems work. Cybercriminals are unable to break down huge combinations into prime components. On the other hand, the Shor algorithm can do this easily, provided it does so on a quantum computer. This algorithm can break asymmetric encryptions that are popularly used today.

The lesson from the Shor algorithm is clear: start preparing today for quantum threats.

What are Quantum Threats

Today, we are using high-performance computer systems to get our work done. However, in technological terms, these are still 'classical' computers from a quantum perspective. The high-performance computing prevalent today will soon evolve into quantum computing. Such computation leverages the frameworks of quantum mechanics to solve complex problems. While classical computers use bits, quantum computing works with qubits. The framework leverages different quantum phenomena to work on multiple possibilities simultaneously.

Complex calculations, the backbone of encryption, can thus be solved quickly.

Why does quantum computing present a cyber risk?

Traditional or classical computers cannot process information quickly. But quantum computers can. This fact means existing encryption methods can be cracked easily. Your organization might pride itself on using robust encryption and following encryption best practices. But, enter quantum computing, and critical digital information won't be safe anymore.

Here are some of the risks of quantum computing, or to be specific, quantum threats:

- **Obsolescence of Public Encryption**

Some of the common types of encryption we use today include RSA (Rivest-Shamir-Adleman), ECC (Elliptic Curve Cryptography), and DH (Diffie-Hellman Key Exchange). The first two are asymmetric types of encryption. The third is a useful fillip to the first to establish secure communication, with secure key exchange. RSA relies on very large integers, and ECC, on very hard-to-solve math problems. As can be imagined, these cannot be solved with traditional computing. However, the time is not far when public-key encryption methods will be rendered obsolete via quantum computing.

- **Harvest Today, Decrypt Later**

Cybercriminals think long-term. They are well aware that quantum computing is still some time away. But that doesn't stop them from stealing encrypted information. Why? They will store it securely until quantum computing becomes readily available; then they will decrypt it. The impending arrival of quantum computers has set the cat amongst the pigeons. Cybercriminals can't wait to score a huge payday later, making concerted efforts today.

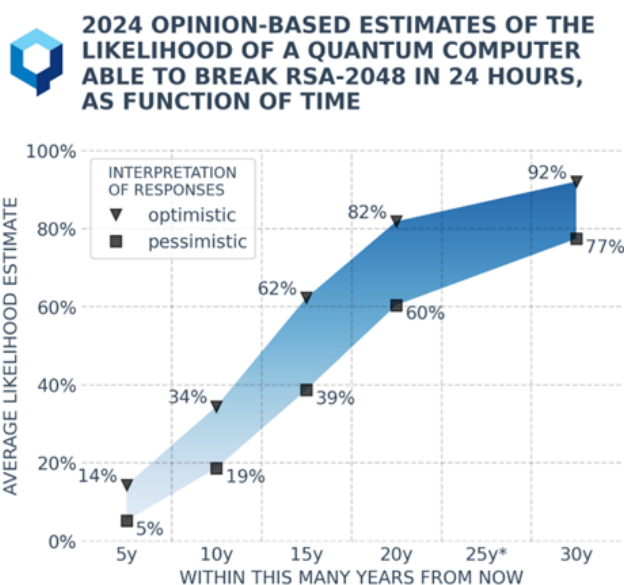
- **Impact on Blockchain**

Blockchain is not unhackable, but it is difficult to hack. A bunch of cryptographic algorithms keep it secure. These include SHA-256 (Secure Hash Algorithm 256-bit) and ECDSA (Elliptic Curve Digital Signature Algorithm). Today, cybercriminals might not attempt to target blockchains and steal crypto. But tomorrow, with the availability of a quantum computer, the crypto vault can be broken into, without trouble.

- **Wide Gap Between Cybersecurity Posture and Attacks**

We keep saying that quantum computing and quantum computing-enabled threats are still some time away. And, this is true. But when the technology is here, it will evolve and gain traction. The attacks will also evolve accordingly. The raw computation power of quantum machines will ensure that encryption can be broken easily. Forging digital signatures, decrypting secure communication, or accessing and decrypting critical information will be a breeze. You therefore need to start upgrading your security systems to future-proof them. Otherwise, you will be caught in a vicious cycle of playing catch-up.

Understanding the Quantum Threat Timeline



The above figure is from a [Quantum Threat Timeline Report 2024](#), by Global Risk Institute. The report pins a timeline for when the Cryptographically Relevant Quantum Computer (CRQC) will be available for

use. It brings together the views of global experts, whose most optimistic estimates suggest the CRQX availability will be between 5 and 30 years.

In this case, optimistic means there is still some time to set up a quantum-resistant infrastructure, but a cautious estimate suggests it can happen within a decade.

There are two critical aspects you must consider in order to plan for a post-quantum threat world.

· **Data Shelf Life**

Take a good, hard look at the data generated by your organization today. Will this data still be relevant or valuable for cybercriminals once the first quantum computer arrives in mainstream use? For example, banking information of a customer or just renewed credit card details (expiry in the next ten years) will be susceptible to quantum threats. This data can be stolen today and decrypted later.

· **Existing Processes**

Do you have plans to update or modify your existing hardware and software that manage and control data? Also, think about the in-development systems and products that will go live. Many of these might become operational when quantum computers are due to arrive on the scene. The fact that these systems and software will be updated over the years means they will be at risk of quantum threats.

Key stakeholders must come together to understand and evaluate the sensitive nature of the data and the systems that process, assess, share, and control this data. Figure out the level of susceptibility these have to quantum threats and devise a remediation plan accordingly.

Why Traditional Encryption Won't Survive

A recent [Google study](#) reported that breaking 2048-bit RSA encryption might only take a week of computation on a quantum machine with 1 million noisy qubits. A slowly dawning reality is that the shift from traditional asymmetric encryption to quantum-resistant encryption must happen now.

Before quantum computing was on the horizon, asymmetric encryption was considered perfectly safe and configured to protect critical data. Existing computing cannot easily factor significantly larger prime numbers that underpin this type of encryption. This is a computationally intensive exercise. It is both time and cost-prohibitive and therefore non-viable.

Quantum computing breaks this non-viability barrier.

The amount of information that quantum computers can store and process is far larger than traditional computers. This means their computational capacity is also significantly faster, enabling them to break asymmetric encryption easily.

We are not far from a scenario wherein a cybercriminal with access to a quantum computer running Shor's algorithm, decrypts RSA and ECC encryption, exposing any private key they guard. This will have repercussions on the everyday security layers that we take for granted and which are delivering immense security value to date. These include [SSL/TLS for web sessions](#), VPN handshakes, PGP and S/MIME

email signatures, OS and app code-signing, and even blockchain wallets or smart contracts.

Until those systems migrate to post-quantum algorithms, they may become redundant soon.

What is Post Quantum Cryptography?

Simply put, post-quantum cryptography helps you address the threats posed by quantum computing. You don't have to worry about Shor's algorithm in conjunction with a quantum computer breaking traditional encryption algorithms. Your quantum-resistant algorithm will keep this threat at bay.

The three categories of quantum-risk resistant solutions include:

- **Post-Quantum Cryptography:** These are new math-based public-key quantum-resistant algorithms.
- **Quantum Key Distribution:** This framework is supported by quantum physics for secure communication channels that share symmetric cryptography keys.
- **Quantum Random Number Generation:** This framework fuels randomness for building robust security protocols.

NIST, or the National Institute of Standards and Technology (NIST), has already come out with a list of post-quantum cryptographic standards primarily covering key exchange and digital signatures. This includes Federal Information Processing Standards (FIPS):

- FIPS 203: This standard, based on ML-KEM (Module-Lattice-Based Key-Encapsulation Mechanism), focuses on general encryption, and is underpinned by small encryption keys that can be shared easily by two parties
- FIPS 204: This standard, based on ML-DSA (Module-Lattice-Based Digital Signature Algorithm), focuses on protecting digital signatures.
- FIPS 205: This standard based on SLH-DSA (Stateless Hash-Based Digital Signature Algorithm), also focuses on digital signatures; the difference being the math approach. If ML-DSA is broken into, this algorithm is the backup.

What does a Quantum Threat Cybersecurity Roadmap look like?

CISOs must start the process to achieve a post-quantum cybersecurity posture today. While different threat timelines offer different estimates, there is no doubt that quantum threats are real and about to arrive.

Here's a roadmap you can use to achieve readiness:

1. Build a Business Case

Make sure you start with preparing a list of tangible impacts to your business with the emergence of post-quantum threats. Also, calculate the ROI of implementing a proactive plan of action before the threats reach your organization's door.

2. Evaluate and Analyze the Threat Perception

Make sure you first understand the at-risk data and processes with respect to quantum computing. This will help you zero in on the post-quantum cryptographic systems that will address a rapidly evolving threat landscape.

3. Asset Prioritization

Create an inventory of the various security protocols used in your organization, including encryption frameworks. Also, create a list of assets that you feel are most susceptible to post-quantum threats.

4. NIST Standards

NIST has already released the list of recommended post-quantum cryptographic standards. Understand these and check how you can begin the implementation process.

5. Begin the Quantum-Safe Journey

Every journey begins with small steps. Methodically start transitioning to a quantum-safe cybersecurity posture. Update your software and hardware, and modify legacy systems. Also, have replacement timelines for systems in place.

6. Continuous Risk Assessment

Your roadmap will evolve as the years go by. New information on post-quantum threats and new solutions will be available. Keep evaluating existing high-risk areas, security protocols, and make incremental improvements.

Final Thoughts

The battle against cyberthreats is won or lost based on your proactive approach (or the lack of it). We are well aware of how organizations dropped the ball on security and experienced an expensive data breach. Do not be that organization. Start taking action and start moving towards post-quantum cryptography to keep your data safe.

About the Author

Ann-Anica Christian is the Content Strategist and Creator at [SSL2BUY](https://www.ssl2buy.com/).

Ann-Anica Christian is an accomplished content strategist and creator with more than seven years of expertise in SaaS, digital eCommerce, and cybersecurity. Her work has been featured on [SSL2BUY](https://www.ssl2buy.com/) Cybersecurity sections, serving as a trusted resource for developers, IT professionals, and business owners. Since beginning her career in technical content development, she has built a reputation for translating complex concepts in website security, IoT, SSL/TLS encryption, and public key infrastructure (PKI) into clear, actionable insights for diverse audiences.

In her current role, Ann-Anica leads content strategies that combine technical depth with business objectives, delivering clarity, accuracy, and engagement across publications. She also authors journals highlighting advancements in technology sectors, reflecting her deep understanding of evolving security threats and best practices, and enabling readers to navigate the fast-changing cybersecurity landscape with confidence.

Connect with Ann-Anica at ann.christian@ssl2buy.com or

our company website <https://www.ssl2buy.com/> to learn more about her work.





CISO of Tomorrow

The CISO of Tomorrow: A Human-Centric Approach to Leading Through AI, Autonomy, and Quantum Uncertainty

The CISO as a Guardian of Trust in a Hyper-Digital World

By Ferris Adi, MBA, CISSP, CISA, CISM, CRISC Cybersecurity Leader & Author, Futurist. CISO Trans Americas Fiber

Picture this: a CEO sits across from you — not to ask about the latest phishing campaign or compliance audit, but to seek guidance on an AI-driven financial model that just made a high-stakes decision without a single human click. Or a board member urgently asks:

"Are we quantum-ready?"

This is the reality of today's CISO.

Gone are the days when firewalls, patch management, and compliance checklists defined the role. The CISO has evolved into a proactive **strategic architect of Trust** in an era where AI acts autonomously, quantum computing looms as an existential cryptographic threat, and digital ecosystems must be resilient, transparent, and ethically governed.

Tomorrow's CISO isn't just a technologist — they're a **futurist, diplomat, and cultural leader**. They must protect innovation without suffocating it, align technology with ethics, and prepare for threats that haven't yet fully materialized.

Here's what that leadership looks like.

1. From Reactive Defence to Autonomous Oversight: Keeping AI in Check

AI has moved beyond automation — it's now an active participant in business decisions. It can approve financial transactions, recommend medical treatments, negotiate contracts, and even execute supply chain purchases.

But with great power comes significant risk: What if an AI makes the wrong call, or worse — is manipulated?

Case Study: The AI Identity Sandbox

A financial institution implemented **short-lived AI identities** — temporary digital credentials that expire within minutes. These AI agents must re-authenticate before performing actions, and every decision is logged, auditable, and tied to strict behavioural rules.

Key safeguards:

- **No rogue AI:** Agents can't act indefinitely or outside their intended scope.
- **Regulatory Trust:** Compliance teams can reconstruct and audit AI-driven decisions.
- **Human-in-the-loop:** AI augments, but doesn't replace, critical oversight.

Lesson for CISOs: Treat AI like a brilliant but unpredictable employee — give it the tools to excel, but ensure guardrails are in place.

2. Zero Trust 2.0: AI Isn't Just a User — It's a Risk Variable

Zero Trust began with the mantra *"Never trust, always verify."* But when AI models constantly evolve, verification must be adaptive.

Case Study: The Adaptive Trust Engine

A European telecom introduced a **dynamic trust scoring system for AI models**. If an AI's behaviour deviates by as little as 3% from expected patterns, its privileges are automatically revoked, and it's sent for retraining or review.

Why this matters:

- **AI drift is real:** Models degrade over time or are subtly manipulated.
- **Real-time detection stops breaches early:** Intervention happens before significant harm.
- **Trust becomes fluid:** Access isn't binary; it's recalculated continuously.

Takeaway: CISOs must expand Zero Trust beyond humans to include the algorithms themselves.

3. Quantum Security: The Cryptographic Doomsday Clock is Ticking

Quantum computing threatens to break much of today's encryption — and adversaries are already **harvesting encrypted data now to decrypt later**. Waiting until "Quantum Day" is a recipe for crisis.

Case Study: Post-Quantum Cryptography (PQC) Rollout

A Canadian government agency adopted a **three-phase PQC migration plan** aligned with NIST's recently finalized standards:

1. **Inventory:** Identify all cryptographic assets and dependencies.
2. **Hybrid Deployment:** Run PQC and traditional encryption side-by-side to ensure backward compatibility.
3. **Agile Orchestration:** Automate cryptographic updates across systems.

Result: They reduced quantum-vulnerable data exposure by 70% in under a year.

Reality check for CISOs: PQC migration is not optional — it's a race against time.

4. Governance Beyond IT: The AI & Quantum Risk Board

AI decisions affect finance, law, healthcare, and public safety. Security leadership must be integrated with ethics, compliance, and strategy.

Case Study: A Healthcare Network's Bold Move

They created an **AI & Quantum Governance Board** — a cross-functional body including legal, compliance, R&D, and security. The board:

- Reviews all AI model change logs.
- Tracks PQC migration milestones.
- Runs annual simulated quantum breach exercises.

The success of the Healthcare Network's approach lies in its cross-functional governance. By bringing together legal, compliance, R&D, and security, the AI & Quantum Governance Board eliminates blind spots and accelerates risk mitigation. This is a model that all organizations should consider.

Lesson: CISOs must speak fluently across law, ethics, and business strategy — not just technology.

5. Culture as the Ultimate Security Control

Even the strongest technical controls fail if an employee unquestioningly trusts an AI's flawed output or ignores a security protocol.

Three human-first cultural practices:

- **Verifiable Explainability:** If AI can't explain its decision in plain language, it doesn't go live.
- **Quantum War Games:** Annual tabletop exercises to simulate post-quantum attack scenarios.
- **Transparency Reports:** Publicly share AI safety posture and PQC readiness metrics.

The cultural truth: Security is not just an IT function — it's a shared organizational value.

Closing Thoughts: The CISO as a Futurist

The CISO of tomorrow is more than a defender — they're a **visionary** who shapes Trust, safeguards innovation, and anticipates threats before they emerge.

- **For CISOs, the time for action is now.** Start piloting PQC migrations, implement AI oversight frameworks, and form cross-functional governance boards. The future of your organization's security depends on the decisions you make today.
- **For Boards & Executives:** Support these initiatives at the strategic level because resilience is now measured in **trust capital** as much as in uptime.

The question isn't whether you'll face AI and quantum threats — it's whether your organization will lead through them or scramble to catch up.

About the Author

Ferris Adi, MBA, CISSP, CISA, CISM, CRISC is CISO at Trans Americas Fiber. He is a Cybersecurity Leader, Strategist, Futurist and Author.

Ferris Adi is an internationally recognized cybersecurity leader with over two decades of experience protecting critical infrastructure, guiding global enterprises, and advising boards on emerging digital risks. His career spans senior security leadership roles across telecommunications, finance, government, and technology—where he has led large-scale security transformations and built high-performing security programs from the ground up.



A trusted voice in the fields of **AI governance**, **Zero Trust architecture**, and **quantum-resilient security**, Ferris blends technical mastery with strategic foresight to help organizations navigate an era where autonomous systems make decisions and quantum computing threatens the foundations of encryption. He is known for translating complex security challenges into practical, business-aligned solutions that build trust while enabling innovation.

Ferris is a sought-after keynote speaker at major conferences and the author of multiple thought leadership pieces on the future of cybersecurity. His work focuses on equipping CISOs and executives with the tools, frameworks, and cultural strategies needed to thrive in the age of AI, autonomy, and quantum uncertainty.

Ferris can be reached on [LinkedIn](https://www.linkedin.com/in/ferrisadi) for collaboration, speaking engagements, and industry discussions <https://www.linkedin.com/in/ferrisadi>



Cyber Insurance: 7 Hard Lessons You Need to Prepare For

By Paul Barbosa, General Manager, Global Cloud Security Business Unit, Check Point

Cyber insurance has matured fast. Insurers scrutinize infrastructure more aggressively, and claims are more complex to process. Most importantly, what was once accepted as the "industry standard" is no longer sufficient.

Your claim will stall if your environment cannot produce the telemetry necessary to reconstruct how an attacker gained access, what they accessed, and how you responded in real-time. Worse, it'll be denied. What matters now is not what tools you say you have in place but whether your system can show evidence of runtime enforcement, detection, and control validation. Here's what that looks like in practice.

1. Your Real Exposure is Being Scanned Without Your Consent

Before you even talk to a broker, the likelihood is that your internet-facing infrastructure is already being analyzed. Underwriters buy data from scanning platforms that identify open ports, expired certificates, known vulnerabilities, and leaked credentials associated with your organization. You might declare multi-factor authentication (MFA) or endpoint detection and response (EDR) in your policy form, but if the external footprint suggests unmanaged services, the risk profile changes.

This external assessment uses DNS enumeration, IP range scans, and breach data correlation. For example, if you have forgotten subdomains pointing to unmaintained servers or left a test S3 bucket open to the public, it impacts underwriting. That alone can drive up your premium or reduce your coverage cap.

2. Policy Enforcement Must be Measurable and Machine-Verifiable

Insurers aren't just asking whether you have controls in place. They are beginning to request logs or reports that prove controls are working. For example, insurers expect MFA enforcement at the identity provider level, role-based access control to be applied and audited, endpoint agents reporting in, and segmentation policies at the network layer.

If you're running infrastructure as code, you need enforcement around that code; simply writing a policy isn't enough. Your configuration management system should alert on drift. Your identity system should raise events when new users are provisioned with elevated privileges. Additionally, you should be able to prove when a misconfiguration was introduced and when it was remediated.

3. Claims Adjusters Rely on Forensic Timelines - You Need Complete and Trusted Logs

When a breach occurs, most organizations enter recovery mode. Yet, insurers go into investigation mode. They will want to know the exact timeline of events, including the initial compromise, lateral movement, persistence techniques, and exfiltration or damage.

That timeline needs to be built from actual logs. If your endpoint logs only store seven days of data, or your SIEM misses certain systems due to licensing constraints, you've already lost context. If cloud audit logs are disabled or not routed to a central location, this issue creates gaps.

4. Ransomware Response Needs to be Automated and Recoverable

Insurers are shifting away from blindly covering ransom payments. They want evidence that you had a response plan, that backups were maintained securely, and that your systems had resilience mechanisms in place. For example, insurers could zoom in on segmentation to isolate affected networks, immutable backups stored out-of-band, and the ability to detect encryption events in progress.

Your claim may be weakened if you lack runtime monitoring capable of flagging mass file changes, encryption behaviors, or beaconing activity. The same applies to backup systems accessible over the same network and sharing the same authentication plane as the compromised systems.

5. Your Cloud Configuration Posture is an Underwriting Risk

Cloud environments evolve fast. Most insurance policies assume that you can keep up with that speed without exposing data or creating escalation paths. If your developers have IAM roles with wildcard permissions or if your default VPCs allow inbound SSH connections, those are surfaced during an incident.

You need real-time config scanning that maps changes to defined baselines. Strategies include marking cloud storage private by default, streaming audit logs centrally, and scoping workload identities to the minimum required privileges.

6. Third-Party Access is Your Liability, Not Theirs

If your vendor is compromised and that compromise leads to data loss, the insurance claim lands on your desk. Your ability to isolate and contain external access is just as important as what you do internally. For example, federated identity, API rate limiting, behavioral monitoring for partners, and the ability to audit what external actors accessed, changed, or initiated.

You must demonstrate that any external integration was properly scoped, permissioned, and documented. For high-risk vendors, you should show real-time access monitoring and alerting, plus the termination of access tokens in the event of a breach.

7. You Will be Audited on Your Detection Coverage

The final stage of any cyber insurance claim is technical review. Your insurer will examine whether your systems were appropriately instrumented to detect what happened. They're looking for coverage across all major vectors: identity, endpoint, network, data, and cloud security.

It's not enough to say you have a detection platform. You need to show:

- It caught something.
- Your team acted on the alerts.
- Escalations happened within defined timeframes.
- Decisions were logged.
- Containment occurred within expected response windows.

Cyber insurance has shifted—do you understand how?

Going forward, organizations that build defensible security architectures—architectures designed not only to prevent attacks but to stand up to post-incident scrutiny—will be best positioned to secure both coverage and claims payouts.

About the Author

Paul Barbosa is the General Manager of Global Cloud Security at Check Point, overseeing Sales, R&D, Marketing, Customer Success, and Engineering. With over 23 years in cybersecurity and communications, Paul brings deep expertise across enterprise, federal, and commercial sectors. A former U.S. Marine and security architect, he's held leadership roles at Skyhigh Security, Cisco, and Cloudflare. His leadership mantra, "Brilliant in the Basics," reflects his commitment to operational excellence and tactical precision. Paul can be reached online at pbarbos@checkpoint.com or <https://www.linkedin.com/in/paulbarbosa/> or at our company website <https://www.checkpoint.com/>





Adapting Incident Response to Autonomous Agents: Evolving Practices for the Future of Cyber Defense

By Tannu Jiwnani, Principal Security Engineer, Microsoft

Introduction

Incident response (IR) is a critical function in cybersecurity, tasked with detecting, managing, and mitigating the impact of security incidents. Traditionally, IR relies heavily on human decision-making, with security analysts investigating alerts, identifying threats, and taking appropriate action. However, as the volume and complexity of cyber threats continue to grow, traditional methods can be slow, error-prone, and insufficient to cope with the increasing pace of attacks.

Autonomous agents AI-powered systems capable of independently detecting, analyzing, and responding to security incidents are rapidly transforming the landscape of incident response. These agents can automate repetitive tasks, enhance threat detection, and execute responses in real time, significantly improving the efficiency and effectiveness of IR. However, the introduction of autonomous agents also necessitates changes in how IR teams operate, collaborate, and make decisions.

This article explores the evolving role of autonomous agents in incident response and outlines how incident response practices should adapt to leverage these new technologies while maintaining human oversight.

1. The Role of Autonomous Agents in Incident Response

A. What are Autonomous Agents?

Autonomous agents in cybersecurity are AI-driven systems capable of acting independently to detect and respond to incidents. These agents can:

- **Monitor network traffic** in real time for signs of malicious activity.
- **Identify and categorize threats** based on machine learning algorithms.
- **Automatically trigger responses**, such as isolating affected systems, blocking malicious IPs, or alerting relevant stakeholders.
- **Analyze incidents** to provide context and recommendations for remediation.

These agents are designed to augment human decision-making, not replace it, by handling repetitive and time-consuming tasks, thus allowing human analysts to focus on more complex or high-stakes decisions.

B. The Need for Autonomous Agents in Incident Response

The increasing sophistication and volume of cyber threats have led to a situation where manual processes are no longer sufficient. Autonomous agents offer several advantages over traditional IR methods:

- **Speed:** Agents can detect and respond to incidents in real-time, reducing the time it takes to contain and mitigate threats.
- **Scalability:** AI-driven agents can handle large volumes of data and alerts simultaneously, allowing IR teams to scale operations without increasing headcount.
- **Accuracy:** By leveraging machine learning, agents can improve the accuracy of threat detection, reduce false positives and ensure that teams focus on genuine threats.
- **Consistency:** Autonomous agents can consistently follow predefined response protocols, ensuring that incident response is not hindered by human error or oversight.

2. How Incident Response Should Evolve with Autonomous Agents

A. Changing the IR Workflow

The introduction of autonomous agents fundamentally changes the workflow of an incident response team. Here's how traditional IR processes should evolve:

Automated Detection and Triage

Traditional IR relies on security analysts to manually review alerts and prioritize them based on severity. With autonomous agents, this process can be automated, as agents can immediately classify incidents, prioritize them based on predefined criteria (such as severity, potential impact, and attack vectors), and

escalate them accordingly. This shift allows IR teams to focus their efforts on high-priority incidents while automated systems handle low-level alerts.

Real-Time Incident Remediation

Historically, remediation efforts in response to an incident have been time-consuming, involving multiple manual steps to isolate affected systems and mitigate the damage. Autonomous agents can carry out initial remediation tasks such as disconnecting infected systems, blocking malicious actors, or initiating network-wide scans automatically. These agents can respond to incidents faster than human analysts, minimizing the window of exposure and reducing the impact of the attack.

Enhanced Collaboration and Communication

Communication across teams during an incident is crucial, but manual coordination can lead to delays. Autonomous agents can automate communication between different stakeholders, including incident responders, IT teams, and management, providing real-time updates and instructions. For example, once an agent detects a threat, it can automatically notify the relevant teams and initiate response procedures, ensuring swift and synchronized action.

Post-Incident Analysis and Reporting

Following an incident, autonomous agents can assist in post-incident analysis by providing detailed reports on the attack timeline, affected systems, and response actions. These agents can also analyze the attack vector and suggest improvements to security protocols, helping teams strengthen their defenses for the future.

B. Maintaining Human Oversight

While autonomous agents significantly enhance incident response capabilities, human oversight remains crucial. Human analysts are still needed to:

- **Provide context:** Autonomous agents can detect and respond to threats, but humans are needed to interpret the broader context, such as understanding business priorities and risk tolerance.
- **Make high-level decisions:** In complex incidents, such as advanced persistent threats (APTs), human decision-making is crucial for determining strategic responses and coordinating with external partners, such as law enforcement or third-party vendors.
- **Refine AI models:** The effectiveness of AI-driven agents depends on continuous training and refinement. Security teams must ensure that their AI models are kept up-to-date with emerging threats and trends.

The key is finding the right balance between automation and human expertise. By automating routine tasks and initial responses, autonomous agents can free up human analysts to focus on more strategic aspects of incident response, such as threat hunting and long-term mitigation planning.

3. Challenges in Adapting Incident Response to Autonomous Agents

A. Integration with Existing Systems

Integrating autonomous agents into existing IR frameworks and tools can be complex. Security Information and Event Management (SIEM) platforms, Incident Response Management tools, and other legacy systems need to be compatible with AI-driven agents to ensure seamless data sharing, communication, and automation. SOCs must carefully plan how to incorporate autonomous agents without disrupting existing workflows.

B. Ensuring Accuracy and Minimizing False Positives

While autonomous agents can significantly reduce response times, they also run the risk of generating false positives or missing critical incidents. Ensuring that agents are properly trained and constantly updated with the latest threat intelligence is essential to maintaining their accuracy. SOCs must continuously monitor agent activity and intervene when necessary to refine their detection capabilities.

C. Security and Privacy Risks

Autonomous agents, if not properly secured, can become a target for attackers themselves. Since these agents have access to sensitive systems and data, they must be rigorously protected against exploitation. Ensuring the integrity and confidentiality of agent-driven actions is crucial to prevent adversaries from compromising the response process.

4. The Future of Incident Response with Autonomous Agents

Looking ahead, the role of autonomous agents in incident response will only continue to grow. Future trends in agent-based IR include:

- **Increased use of AI-driven agents** for proactive threat hunting and advanced incident detection.
- **More sophisticated integration** with machine learning models that improve incident prediction and response time.
- **Full integration** with threat intelligence platforms, enabling agents to respond to emerging threats faster and with greater precision.

As the capabilities of AI and machine learning continue to evolve, autonomous agents will play an even greater role in shaping the future of incident response, making it faster, more efficient, and better suited to the complex and dynamic nature of modern cyber threats.

Conclusion

The introduction of autonomous agents into incident response workflows represents a paradigm shift in how SOC's approach threat detection, management, and remediation. By automating repetitive tasks, accelerating response times, and improving the accuracy of threat detection, agents significantly enhance the effectiveness of incident response teams. However, to fully realize the potential of autonomous agents, SOC's must adapt their processes and integrate these agents, while maintaining human oversight and refining AI models to minimize false positives.

As cybersecurity threats continue to evolve, the integration of autonomous agents into incident response is no longer optional it is a necessary step toward achieving greater resilience and agility in defending against cyberattacks. With the right combination of automation and human expertise, SOC's can stay ahead of emerging threats and respond to incidents faster and more effectively than ever before.

About the Author

Tannu Jiwnani is a Principal Security Engineer at Microsoft with over a decade of experience in cybersecurity, specializing in cloud security, incident response, and machine learning integration in security tools. Tannu has a deep focus on building scalable, resilient security systems and improving incident response frameworks within organizations. With a strong background in Identity & Access Management (IAM) and threat actor detection, Tannu is passionate about enhancing organizational defenses against evolving cyber threats and can be reach out at [Linkedin](#)





Adversarial GRC Weaponizing Compliance to Defend against Synthetic Threats

By Victor D Patterson Sr, AI Cybersecurity Strategist & Founder of DeepSecure™ — Architect of AI Security Doctrine

Traditional Governance, Risk, and Compliance (GRC) frameworks often function as passive systems, focused largely on meeting regulatory requirements after issues arise. Adversarial GRC transforms this model into an active cyber defense mechanism, leveraging predictive strategies, adversarial intelligence, and proactive compliance enforcement. By weaponizing compliance, organizations can preemptively counter emerging threats, particularly those leveraging sophisticated techniques like AI-driven cyberattacks and deepfakes.

Defining Adversarial GRC

Adversarial GRC applies cyberwarfare methodologies to compliance management, integrating:

Proactive Risk Exploitation: Proactively identify regulatory gaps and compliance weaknesses before attackers can exploit them, simulating real-world scenarios to anticipate and mitigate threats.

OSINT Reconnaissance: Use open-source intelligence (OSINT) and frameworks such as MITRE ATT&CK to map compliance vulnerabilities and assess policy robustness.

AI-Driven Threat Modeling: Utilize artificial intelligence to predict potential compliance failures, such as exploiting gaps in identity verification processes or regulatory blind spots.

Cyber Deception: Employ deceptive tactics like synthetic adversary simulations to continuously test and enhance internal controls against deepfake and fraud attacks.

Weaponized Compliance: Strengthen and proactively exceed existing regulatory standards to preemptively counter potential threats.

By actively engaging in these practices, compliance becomes a dynamic defense rather than a static checklist.

Strategic Cybernetics and Compliance as Warfare

Adversarial GRC is deeply influenced by strategic cybernetics, a legacy of Norbert Wiener, emphasizing adaptive feedback loops and anticipatory defense. It reconceives compliance as code—continuously refined and adjusted to evolving threats. Organizations adopting this mindset treat cybersecurity governance akin to military strategy, embedding resilience, adaptability, and proactive engagement into their compliance and risk management frameworks.

Comparison with Existing Models

Zero Trust Architecture (ZTA)

Zero Trust principles, summarized as "never trust, always verify," complement Adversarial GRC by reinforcing proactive threat management. However, while Zero Trust primarily addresses technical and access control mechanisms, Adversarial GRC broadens this lens to governance policies and procedures, directly challenging policy robustness against adversarial tactics.

AI-Driven Security Operations (AI-SOCs)

AI-SOCs leverage artificial intelligence for real-time threat detection and incident response. In contrast, Adversarial GRC operates at the governance layer, focusing on proactive, policy-driven threat mitigation and compliance resilience, essentially creating a synergy between operational defense and strategic governance.

GRC Automation and Continuous Compliance

Automated GRC processes streamline compliance but typically focus on efficiency. Adversarial GRC enhances this by embedding adversarial thinking, continuously testing and adapting compliance frameworks to anticipate and counteract innovative threats proactively.

Alignment with Academic and Governmental Approaches

While still emerging, Adversarial GRC builds on well-established research in active defense, adversarial thinking, and cyber resiliency. Notably, government strategies, including the U.S. Department of Defense's Zero Trust Strategy, increasingly emphasize proactive, anticipatory cyber defense, closely aligning with the principles of Adversarial GRC.

Programs such as the Cybersecurity Maturity Model Certification (CMMC) demonstrate government recognition of compliance as a proactive security measure. Adversarial GRC extends these concepts further by treating compliance not merely as regulatory adherence but as a strategic weapon against cyber threats.

Industry Adoption and Future Potential

Early adoption of Adversarial GRC concepts can be observed in organizations integrating red-team insights into governance frameworks and utilizing threat intelligence to inform compliance strategy. With emerging trends in AI-driven compliance management and predictive threat modeling, industry interest is poised to expand significantly.

Adversarial GRC also faces challenges, notably the need for standardized frameworks, skill development, and cultural shifts within organizations traditionally separating compliance and cybersecurity operations. However, the growing threat landscape, driven by AI-powered adversaries and evolving attack vectors, underscores the urgent necessity for integrating adversarial strategies into GRC practices.

Gaps and Challenges

Key challenges to widespread adoption include:

Lack of standardized frameworks and metrics to measure efficacy.

Cultural shifts required for compliance professionals to adopt adversarial and proactive mindsets.

Tooling gaps for automating and systematically implementing adversarial strategies within compliance.

Ensuring that aggressive compliance strategies remain within legal and ethical bounds.

Overcoming these barriers will require collaboration between academia, industry, and regulatory bodies to develop standardized methodologies and training programs that merge adversarial thinking with governance practices.

Conclusion: Toward a New Doctrine

Adversarial GRC represents a significant evolution in cybersecurity governance, transforming compliance from reactive checklists into proactive defense mechanisms capable of anticipating and neutralizing threats. By adopting this strategic mindset, organizations can significantly enhance their resilience against sophisticated synthetic threats, securing their digital ecosystems against the innovative attacks of tomorrow.

As organizations, governments, and industry bodies recognize the critical importance of proactive governance, Adversarial GRC is positioned to become a foundational cybersecurity doctrine, shaping how compliance and risk management are conducted for decades to come.

Author Bio: Victor D. Patterson Sr. is a cybersecurity strategist, reverse engineer, and researcher specializing in adversarial AI defense and strategic governance frameworks. His pioneering work in Adversarial GRC has positioned him as a thought leader at the intersection of compliance, cyber warfare, and artificial intelligence.

References & Strategic Citations

This article draws inspiration and strategic alignment from the following foundational frameworks and publications:

- **MITRE ATT&CK® Framework**

For adversarial modeling, OSINT-based reconnaissance, and mapping of cyber threat behaviors.
<https://attack.mitre.org>

- **NIST SP 800-53 & SP 800-172**

For comprehensive compliance control baselines and cyber-resilience through enhanced security requirements.

<https://csrc.nist.gov/publications>

- **U.S. Department of Defense Zero Trust Strategy (2022)**

For anticipatory cyber defense, continuous verification, and trust segmentation.
<https://dodcio.defense.gov/Portals/0/Documents/Library/DoD-Zero-Trust-Strategy.pdf>

Citation & ORCID Registry

Victor D. Patterson Sr. is officially registered with **ORCID**:

 <https://orcid.org/0009-0005-5258-3548>

This article is also documented and preserved via Zenodo under Creative Commons License:

 <https://zenodo.org/records/15239139>

APA Citation:

Patterson, V. D. (2025). *The Age of Deepfakes™ – A GRC Strategy Doctrine for Protecting Digital Ecosystems (Featuring DeepSecure™, Cerberus Shield™, Odin's Code™, and GRC-AIDE™)* (v1.0 – Public Preprint / Strategic Doctrine, Pending IEEE Submission). Zenodo. <https://doi.org/10.5281/zenodo.15239139>

License: CC BY-ND 4.0 (Attribution–NoDerivatives)

About the Author

Victor D Patterson Sr is the Founder & AI Security Architect | DeepSecure™ | ISACA 2025 Speaker | RTC/IEEE 2025 Presenter, & Speaker of the DeepSecure™ (Framework Author)]. He Victor is the original creator of **DeepSecure™**, the world's first AI-powered SOC doctrine built to detect and counter AI-enabled threats such as deepfake deception, synthetic identity warfare, and adversarial AI fraud. His frameworks are validated against leading standards such as **NIST AI 600-1**, **ISO/IEC 42001**, and **MITRE ATT&CK**.



He is also the architect of **Cerberus Shield™**, a layered GRC defense model combining adversarial AI deception modeling, post-quantum compliance, and predictive risk intelligence. His third doctrine, **Odin's Code™**, fuses cyberwarfare strategy with AI governance and adversarial reconnaissance.


Victor's published research, *The Age of Deepfakes: A GRC Strategy Doctrine for Protecting Digital Ecosystems*, is archived on Zenodo and recognized by ISACA and RTC/IEEE as a future-forward cyber governance model.

"I didn't design for quantum compliance—I built a system that forced it into existence."

Victor is a rising voice in adversarial GRC, AI-driven SOC automation, and national-scale cyber defense strategy. His independent doctrine has earned recognition from cybersecurity leaders, academic institutions, and public sector collaborators across North America.

As an ISACA 2025 speaker and RTC/IEEE 2025 presenter & speaker at Illinois Institute of Technology, Victor continues to push cybersecurity into its next era—AI-driven, adversarially hardened, and globally resilient.

Victor can be reached online at vpatterson87@gmail.com | VictorDeepSecure.ai@proton.me | <https://www.linkedin.com/in/victor-patterson/> and at our company website DeepSecure.ai



The Future of Mobile App Security: Emerging Technologies and Trends

By Jason Cortlund, Mobile App Security Evangelist at Guardsquare

The mobile threat landscape is evolving fast, and it's shaped by API abuse, runtime threats, [AI-powered social engineering](#), and more. Mobile app security used to be about passing app store checks and applying baseline protections. Now, it's about managing real, evolving threats across the entire app lifecycle. Meanwhile, the pressure to move faster, meet compliance standards, and deliver frictionless experiences hasn't let up.

But many organizations are still playing defense. [New research](#) conducted by Enterprise Strategy Group (ESG) reveals a growing gap between perceived protection and real-world risk. While 93% of organizations believe their [mobile app security](#) is sufficient, the same study shows they're facing an average of 15 security incidents per year, with each incident costing nearly \$7 million.

Even still, 85% of organizations prioritize security after an incident occurs. This approach is risky, given the rising stakes and consumers' increased awareness of mobile app security issues than ever before. However, as threats become more pervasive, a proactive approach to mobile app security is quickly increasing. Here are some of the biggest trends to watch that will help improve your overall mobile app security posture.

Key Trends Shaping Mobile App Security

Zero Trust Moves to Mobile

Zero Trust Architecture (ZTA) is no longer confined to enterprise networks. On mobile, this means that every access request — regardless of the user, device, or location — must be verified in real-time.

Unlike perimeter-based models, [ZTA assumes every connection is untrusted](#) until proven otherwise. When applied to mobile app security, this approach limits exposure, even if an attacker compromises a device. The result is tighter, context-aware access control. This is a critical defense, as mobile apps increasingly handle sensitive data in banking, healthcare, identity verification, and other areas.

Runtime Protection Becomes the New Baseline

As attackers get more sophisticated, static mobile app protections like [obfuscation](#) need to be combined with [Runtime Application Self-Protection \(RASP\)](#). RASP brings mobile apps real-time awareness of their operating environment, detecting mobile app security threats like code injection, hooking frameworks, or rooted environments at runtime.

They give mobile apps real-time situational awareness, enabling defensive actions such as shutting down the app, restricting functionality, or alerting the security team. Expect RASP adoption to accelerate as businesses look for deeper visibility and resilience on untrusted endpoints.

Secure SDLC Gains Ground

According to the Enterprise Security Group's findings referenced above, 74% of organizations report that their app development teams are under increased pressure to move faster. In comparison, 71% indicate that this speed pressure has compromised mobile app security. Many teams are becoming aware of this duality and are addressing it by shifting security left. Instead of bolting on mobile app protections post-release, more teams are [embedding security directly into the software development life cycle \(SDLC\)](#) — from requirements gathering to testing and deployment.

This secure SDLC model reduces long-term costs, surfaces risks earlier, and creates closer alignment between engineering and security teams. It also aligns well with continuous delivery models, enabling faster iteration without compromising protection.

Mobile APIs Under Attack

[Mobile APIs](#) are a growing target. Attackers are exploiting poorly protected endpoints to extract data, manipulate app behavior, or impersonate users. In fact, mobile API abuse has already led to real-world breaches, especially in industries handling payments, healthcare records, or PII. For example, in 2024, a multi-factor authentication app [Authy, experienced an API endpoint breach](#), in which attackers accessed and published millions of Twilio users' phone numbers.

[Securing mobile APIs](#) now requires more than rate limiting. Development teams need to layer in defenses like [mobile app attestation](#) and token binding to ensure only untampered, legitimate apps can access backend APIs, to ensure it's *your* app interacting with *your* APIs. This step helps block impersonation attempts and API scraping, both of which are rising among credential stuffing and bot-based attacks.

On-Device Privacy and Threat Detection

Cloud-based monitoring is useful, but [on-device intelligence](#) is gaining traction — and for good reason. Processing threat signals directly on the device enables faster response times and better privacy controls. With on-device [mobile app threat detection](#) and attestation, apps can verify the integrity of their environment and make decisions in real time — you don't need to upload sensitive user data to the cloud to spot abnormal behavior.

Instead, you can detect jailbreaks, hook attempts, or suspicious message patterns locally and act immediately. This approach improves security while aligning with evolving data privacy regulations that restrict data transfer and storage.

Regulation Tightens the Screws

Around the world, [compliance mandates](#) are becoming more prescriptive. Frameworks like GDPR, CPRA, and PCI DSS now require mobile apps to enforce encryption, limit data collection, and conduct regular security audits.

These regulations are forcing mobile app security into product strategy conversations earlier. For global brands, adapting to local and international compliance requirements will quickly become table stakes.

What a Proactive Mobile App Security Strategy Looks Like

To address these trends and more, organizations are embracing [multi-layered mobile app security strategies](#). These include a combination of techniques, such as:

- [Code hardening](#) and **encryption** to resist reverse engineering and protect IP
- **Runtime protection** to detect tampering, debugging, and dynamic attacks
- [Mobile app security testing \(MAST\)](#) to uncover issues in code and third-party SDKs
- **Real-time threat monitoring and attestation** to surface real-world attack behavior and unauthorized API access, guiding the response

Case in point: A [top Central American bank](#) transitioned away from a low-support, cloud-wrapped security vendor after experiencing crashes and limitations. With a multi-layered approach — including code hardening, testing, and real-time threat monitoring — the bank improved stability, passed [pentesting](#), and now actively tracks threats in production.

These methods are most effective when integrated directly into development workflows. In fact, 46% of organizations surveyed in "The Growing Threat Landscape" say that developer-friendly security tools are a top priority. Nearly 60% plan to increase security budgets, with ease of use and automation among the most significant drivers.

Looking Ahead

Security must keep pace with innovation. As AI changes how threats are delivered and detected, and as regulations tighten, mobile app security will remain as fast-moving as ever. The good news? Organizations are becoming increasingly aware of these threats and are prepared to take action.

To respond, organizations must build with mobile app security in mind from day one. By integrating protection, testing, and monitoring throughout the mobile app development lifecycle, teams can reduce risk, improve resilience, and protect both users and their bottom line.


Forward-looking security teams are already adapting. They're investing in mobile app protection that runs where the risk lives: on the device, in real time, with tools that developers can actually use.

About the Author

Jason Cortlund is the Mobile App Security Evangelist at [Guardsquare](https://www.guardsquare.com)

More information can be found at <https://www.guardsquare.com>





More Than Just a Harmless Box: Making the Case for Printer Platform Security

By Steve Inch, Global Senior Print Security Strategist & Product Management Lead at HP Inc.

These days, CISOs and their teams have their work cut out for them. It is their job to identify and reduce security risks, as well as monitor for signs of intrusion across an ever-growing attack surface. Against this backdrop, they are also under pressure to deliver cyber resilience and meet increasing regulatory demands, both at a national and industry level.

But for many organizations, there is a major gap in their capabilities: printer security. Too often, printers are overlooked as a source of cybersecurity risk. For example, researchers identified [half a million internet-connected printers](#) that were vulnerable to hijacking.

Although IT teams spend nearly four hours a month managing hardware or firmware security, they frequently lack the visibility and control required to mitigate security risks across their print fleets. With the average print-related data loss costing organizations £820,000, according to the Ouocirca Print Security Landscape 2025 report, it's clear that printer security is an area that enterprises can't afford to ignore and that managing printer security well should be part of every CISO's plan to improve cyber resilience.

Why printer security matters

Whereas PCs typically benefit from greater security visibility and more mature endpoint protection measures, printers are often overlooked and may lack adequate defenses altogether. This oversight leaves a critical gap in the organization's security posture. The truth is that today's corporate printers are far from "harmless boxes" — they're sophisticated, networked devices with onboard storage and internet connectivity, making them an attractive and viable target for threat actors.

By exploiting firmware vulnerabilities, manipulating hardware components, and taking advantage of poor configuration (such as open ports and default credentials), malicious actors could target internal storage for sensitive documents. Or they could compromise a printer to gain entry to the corporate network for more serious data exfiltration and digital extortion. They could even hijack machines to conscript them into a botnet, such as Mirai, to launch attacks on other targets, including public websites, third-party networks, or other unsuspecting organizations.

Securing the lifecycle

There's risk at every stage of the printer lifecycle. [Research from HP Wolf Security reveals](#) that during the procurement process, IT and security decision makers (ITSDMs) are often excluded from assessing vendor security claims. And fewer than 40% of global firms bring IT, security, and procurement together to define security standards when purchasing printers. This means that many fail to request key technical documentation to validate security claims, or to require printer manufacturers to respond to security-related questions.

Security challenges continue when the printer lands on IT's desk. Around half of ITSDMs can't confirm if a printer has been tampered with, whether in the factory or in transit. And then there's ongoing management of printer security. Just 36% of global organizations apply firmware updates promptly despite claiming to spend an average of 3.5 hours per printer per month on hardware or firmware security management.

Worryingly, many respondents say they have difficulty identifying vulnerable printers, tracking unauthorized hardware changes, and ensuring firmware/BIOS compliance with IT security policies. They also struggle with detecting security events linked to hardware- and firmware-level attacks. Beyond cyber threats, 70% of ITSDMs are also increasingly concerned about offline threats, such as people printing and taking away sensitive company information.

The challenges continue right to the end of the lifecycle. Many ITSDMs lack confidence in current sanitization solutions, with 35% uncertain whether printers can be fully and safely wiped. This uncertainty creates significant roadblocks to sustainable disposal, with 86% citing data security as an obstacle to reuse, resale or recycling — and 39% describing it as a "major" or "severe" concern. A quarter of global organizations believe they must physically destroy printer storage drives to mitigate data risks, while 1-in-10 go even further, insisting on destroying both the device and its storage. As a result, otherwise usable devices are taken out of circulation, depriving organizations of potential extra revenue and undermining their sustainability efforts.

Cyber resilience starts here

To plug these gaps, organizations need to prioritize improvements across the printer lifecycle. To start with, IT, security and procurement teams must collaborate more effectively to define security and resilience requirements for new printers. Involving ITSDMs earlier in the procurement process also helps ensure the right questions are asked and vendor claims properly verified.

Next, the ongoing management of print security deserves much greater focus, including managing firmware admin passwords to securely enable configuration changes and technical support. Control of changes to physical components, along with regular checks of device security configurations, is also essential for maintaining printer security. When managing printer fleets, prompt application of firmware updates and using Security Information and Event Management tools to monitor security logs from the printer fleet are a must to reduce the attack surface and detect exploitation attempts.

It's vital IT teams look to buy printers that can be secured across the whole lifecycle. Organizations require models that can continuously monitor for zero-day threats and malware, and those that support secure encrypted printing and data loss prevention. Machines with built-in secure erasure of hardware and firmware data will help to streamline second life and recycling. By ensuring any printer supports multi-pass overwriting of sensitive print data on the hard disk drive and the secure erase of solid-state drives, organizations can gain peace of mind that private data can't be retrieved after decommissioning.

A long-term commitment

Despite the proliferation of digital workflows, printers continue to be an integral part of the modern office. But as such, they are also an attractive target for threat actors and a potential cybersecurity risk. Organizations committed to cyber resilience must not overlook their responsibility to mitigate this risk as comprehensively as possible.

Above all, remember that printer security is a long-term commitment, with refresh cycles spanning years. It's time to give printer security the attention it deserves.

About the Author

Steve Inch is Global Senior Print Security Strategist and Product Management Lead at HP Inc.

Steve can be reached online at LinkedIn:

<https://www.linkedin.com/in/inchsteve> and at our company website <https://www.hp.com/>.





Preparing for Updated HIPAA Security Requirements

By Erik Eisen, CEO of CTI Technical Services

Whether it's implemented as proposed or with modifications based on stakeholder comments, the HIPAA Security Rule will be overhauled by 2026—the first major update since 2013.

Cyberattacks have reached unprecedented levels in both volume and sophistication since the HIPAA Security Rule was first implemented 20 years ago. In 2013, healthcare organizations experienced just 269 data breaches compared to 725 in 2023. According to the Office of Civil Rights (OCR), hacking-related data breaches against healthcare organizations increased 239% between January 1, 2018, and September 30, 2023, while ransomware attacks increased 278% over the same period. The largest healthcare attack in history occurred just one year ago when approximately 190 million individuals, roughly 57% of the U.S. population, were affected by the Change Healthcare breach.

With cybersecurity risks at an all-time high, few in the industry question the need to modernize the regulations. However, even if OCR modifies the final rule to address concerns expressed during the public comment period, compliance will be a heavy lift for many healthcare organizations. That reality, coupled with the common-sense need for robust security around protected health information (PHI) and other patient data, means healthcare organizations must take steps toward compliance now.

The Evolution of HIPAA Security

In the executive summary of its HIPAA Security Rule to Strengthen the Cybersecurity of Electronic Protected Health Information proposed rule, published to the Federal Register in January 2025, OCR was clear about the catalyst for change, writing that technology has reshaped the healthcare industry and the environment in which care is provided.

“Cybersecurity is a concern that touches nearly every facet of modern healthcare, certainly more so than it did in 2003 or even 2013. Almost every stage of modern healthcare relies on stable and secure computer and network technologies ... Thus, cyberattacks, malfunctions, and inadvertent errors can negatively affect the provision of health care, as well as the efficiency and effectiveness of the health care system,” writes the OCR.

Adding to healthcare’s risk profile are covered entities (CEs) and business associates (BAs), where unintentional and nefarious events can endanger electronic protected health information (ePHI) and other sensitive data. Thus, OCR determined that it was time to update the rule to address:

- Technology advancement.
- Shifting trends in breaches and cyberattacks.
- OCR’s greater enforcement experience.
- Improvements in guidelines, best practices, methodologies, procedures, and processes for protecting ePHI.
- Legal decisions impacting Security Rule enforcement.

The goal is also to re-address one of OCR’s most significant challenges when it comes to regulating security: the rapid advancement of both health IT and the methods employed by malicious actors.

Too-prescriptive mandates would necessitate updating the rule—an onerous, costly, and time-consuming process—more frequently than is realistic. Previous iterations of the HIPAA Security Rule attempted to address this by being flexible with compliance. Many security measures were also classified as “addressable implementations,” meaning they were strongly recommended but not explicitly required.

For example, the current rule requires any organization touching ePHI to conduct a security risk assessment to evaluate potential risks and vulnerabilities, resolve any identified vulnerabilities, and document the steps taken. OCR even provides a tool for use in conducting the evaluation. But beyond that, there is no prescriptive guidance. As a result, many healthcare organizations that lacked the resources or technical knowledge to conduct a comprehensive risk assessment wound up taking shortcuts.

The lack of prescriptive guidance, coupled with limited verification requirements, led to unintended consequences, including non-compliance. The increased specificity and expanded requirements in the proposed rule should close those loopholes and harden security around health information.

A Heavy Lift

Industry support for the HIPAA Security Rule overhaul is broad, as are concerns that the compliance burden will be too high for many organizations affected by it. There was a consensus throughout the nearly 4,750 letters submitted to OCR during the proposed rule's two-month public comment period that many requirements would be nearly impossible for some organizations to meet without resource assistance. These include:

- Documenting all Security Rule policies, procedures, plans, and analyses.
- Developing and maintaining a technology asset inventory and network map of ePHI movement.
- Identifying potential vulnerabilities and predisposed conditions.
- Notifying changes to or termination of an employee's ePHI or information system access within 24 hours.
- Establishing written procedures to restore the loss of specific information systems and data within 72 hours.
- Conducting annual compliance audits and regular risk assessments.
- Collecting annual written verification of compliance from BAs.
- Encrypting ePHI when it is at rest and in transit.
- Establishing technical controls for consistent configuration of information systems.
- Use multifactor authentication (MFA).

Additionally, the proposed rule converts many addressable implementation specifications to required, which eliminates a core flexibility aspect of the rule. Finally, for many, compliance with the updated HIPAA Security Rule will not be feasible with their existing technical infrastructure. It would necessitate significant investments in new technologies capable of protecting ePHI as mandated by the rule.

Lessening the Burden

The good news is that compliance does not have to come at the cost of financial ruin. Small steps toward anticipated mandates can be taken now—many of which are common-sense protective measures that should be deployed regardless of the requirements in the final rule.

For example, even organizations with limited budgets can implement MFA, which is a highly effective yet reasonably priced protection against phishing and other forms of infiltration. Regularly backing up data now will ensure continuous access to information in the event of a system outage. At the same time, ransomware or exfiltration protection that goes beyond encryption can prevent bad actors from exploiting vulnerable access points once they are inside a system.

Other actions healthcare organizations can (and should) take now include conducting a security risk assessment and drafting a mitigation and remediation plan. Doing so allows for the prioritization of limited resources.

It is also likely that even well-resourced healthcare organizations will find it infeasible to implement these early steps or achieve compliance within the timeframes outlined in the final security rule without third-

party support. Thus, now is the time to identify the right trusted IT management firm to assist with enhanced security and, eventually, regulatory compliance.

Look for firms with a deep understanding of healthcare-specific compliance requirements. Prospective partners should also offer comprehensive services to ensure they can address the comprehensive needs related to compliance with the HIPAA Security Rule and other issues that may arise, including the ability to future-proof security. They should also possess advanced expertise and the willingness and ability to leverage cutting-edge tools and processes that can outperform older or less adaptive technologies.

Look for a partner that emphasizes long-term relationships and offers personalized customer support. Other must-haves include flexibility and scale in their approach to services, transparent price structures, and simple contracts with clear and fair service terms. Finally, during the evaluation process, be sure to ask prospects about response times and disaster recovery capabilities and obtain—and check—references.

A Proactive Approach

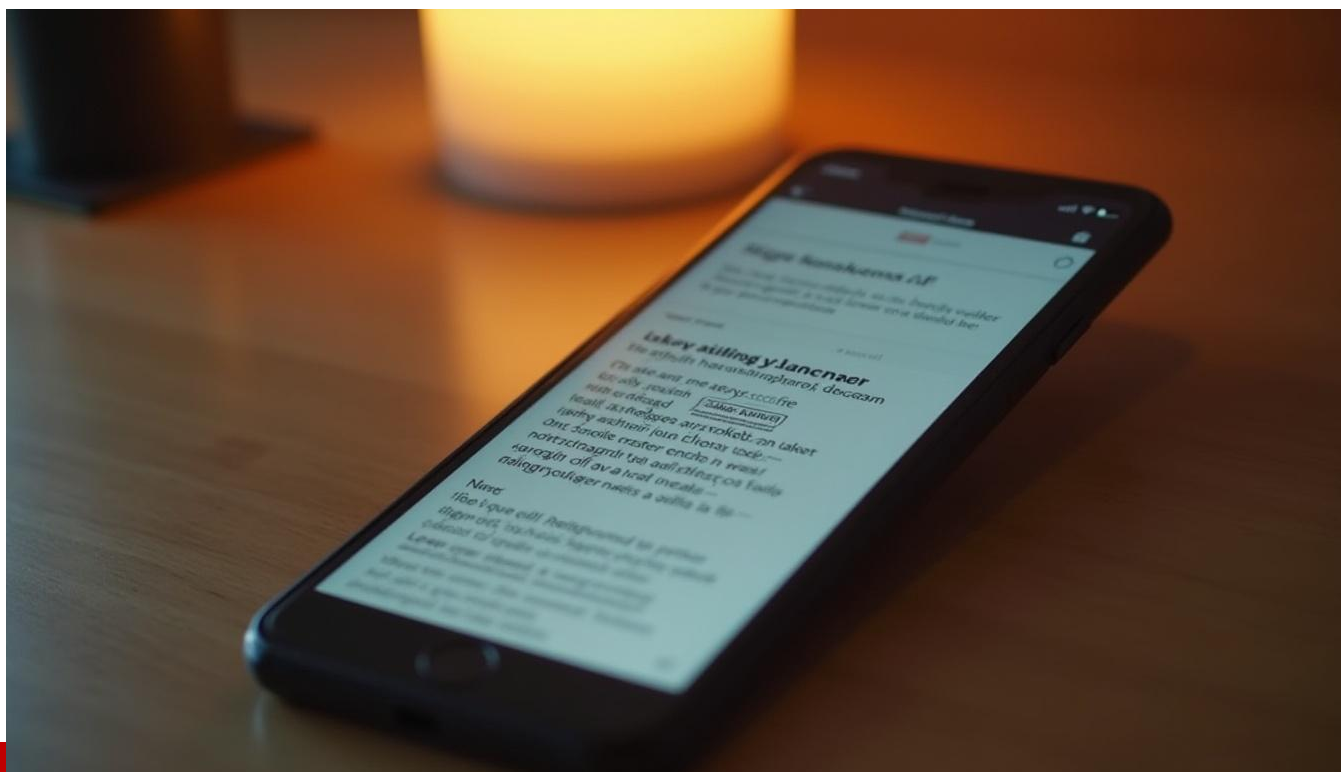
While the final requirements may differ from what has been proposed, there is little likelihood that OCR will retract its decision to overhaul the HIPAA Security Rule. It is an action that is long overdue and should serve as a reminder that strengthening data protection is the right thing to do, whether mandated by OCR or not.

Taking steps now to prepare for the inevitable will significantly ease compliance burdens and improve the protection of one of healthcare's most valuable assets. For those hospitals, health systems, physician practices, and other impacted healthcare organizations with limited resources, identifying the right IT management partner now and taking small steps toward compliance will put them on the right path to protecting patient data.

About the Author

Erik Eisen is the CEO of [CTI Technical Services](#), a leading provider of IT support and cybersecurity services, serving a diverse clientele across various industries, including healthcare, dental, hospitality, legal, manufacturing, and others. For more than 20 years, Erik has provided security and cybersecurity, implemented state-of-the-art technology solutions, and delivered services that protect the integrity of businesses' data and, more importantly, their clients' and customers' data. He is at the forefront of exploring AI integration to help enterprises enhance operations while preserving the essential human touch, ensuring that new technologies are effectively adopted to improve client service.





Email Threats Are Evolving: Why Low-Tech, AI-Powered Attacks Are Outpacing Traditional Defenses

By Usman Choudhary, Chief Product and Technology Officer, VIPRE Security Group

The cyber battlefield is shifting, and attackers are gaining the upper hand by taking a deceptively simple approach. While many security professionals are on high alert for zero-day exploits and advanced persistent threats, the most successful email attacks in 2025 are often low-tech in nature, exploiting human behavior over system vulnerabilities. Our recent analysis at VIPRE Security Group, which examined more than 1.45 billion emails across Q1 and Q2 of this year, reveals that cybercriminals are increasingly using artificial intelligence to scale and personalize their campaigns, targeting people rather than systems with alarming success.

The findings show a stark shift in email-based threats. Phishing, malware delivery, and business email compromise (BEC) attacks are evolving beyond link-based or easily detected vectors. Instead, they now rely on hyper-personalized deception and social engineering techniques that legacy security tools struggle to identify.

Phishing kits go dark

In Q2 2025, 58 percent of phishing kits detected were unidentifiable. This marks a significant departure from previous years when threat researchers could often track specific kits and campaigns. The rise of custom-built or heavily obfuscated phishing kits suggests that threat actors are avoiding known signatures and using AI to create scalable, low-detection phishing infrastructure.

Named kits such as Evilginx (20 percent), Tycoon 2FA (10 percent), and 16shop (7 percent) remain present, but the majority of kits now defy categorization. These kits are often deployed across global infrastructure, making detection and attribution increasingly difficult.

Callback phishing gains traction

One of the most surprising trends this year is the emergence of callback phishing. In Q1 2025, this tactic accounted for 16 percent of phishing attempts, up from almost nothing in 2024.

In callback phishing, attackers send emails that prompt users to call a phone number to resolve an issue. These emails often appear to be from trusted sources like IT departments or service providers. Once the victim calls, they are socially engineered into disclosing credentials or installing malware. Because no links are included in the initial email, traditional phishing detection methods often fail to flag these messages.

SVG files: a rising delivery mechanism

Although PDF attachments remain the most common file format used in phishing campaigns, accounting for 36 percent of cases, SVG files are closing the gap quickly at 34 percent. SVGs, or scalable vector graphics files, are attractive to attackers because they can include embedded JavaScript code through the use of script tags. When opened in a browser, these scripts can redirect the user to a malicious site or initiate a file download.

Many security filters do not scan SVGs as aggressively as they do other file types, making this a favored tool for bypassing defenses. The United States remains the most targeted region for this type of attack, followed by countries in Western Europe.

Malware-as-a-Service matures

In both quarters analyzed, Malware-as-a-Service (MaaS) played a key role in the spread of dangerous payloads. In Q1, the backdoor malware XRed accounted for three times more infections than any other malware family. In Q2, Lumma Stealer took the top spot. Both malware families are sold on underground forums as MaaS, making them accessible to a wide range of cybercriminals, including those with limited technical expertise.

These malware variants are commonly delivered via malicious .docx, .html, or .pdf files hosted on legitimate cloud services such as Google Drive or Microsoft OneDrive, which further complicates detection efforts.

Industry-specific targeting intensifies

Manufacturing continues to be the most targeted sector in the email threat landscape. For six consecutive quarters, manufacturers have been the top target for cybercriminals. In Q2 2025, the industry accounted for 26 percent of all email-based attacks. In Q1, that number was even higher at 36 percent.

Manufacturers are particularly vulnerable due to their reliance on email communications across supply chains, vendor relationships, and logistics. Retail and healthcare followed in Q2 with 20 percent and 19 percent of attacks, respectively. In Q1, the financial sector tied with retail for second place at 15 percent each.

Business email compromise becomes more local

BEC attacks are becoming more linguistically sophisticated. In Q2, 42 percent of impersonation attempts targeted English-speaking executives, while 38 percent were aimed at Danish-speaking leaders. Swedish and Norwegian-speaking executives made up an additional 19 percent of targets.

Attackers are now localizing their messages, using native-language communications to increase the credibility of their impersonation attempts. CEOs and senior executives remain the most common targets, comprising 82 percent of all impersonation efforts. However, other departments are also in the crosshairs, including directors, human resources, IT staff, and even education administrators.

New delivery mechanisms and tactics

Phishing links are increasingly delivered through open redirect mechanisms. In Q2, 54 percent of malicious links were hidden behind open redirects on legitimate services such as email marketing platforms and analytics tools. These links obscure the final destination, making them harder for users and filters to identify as suspicious.

There is also a growing use of QR codes embedded in PDF attachments. These codes entice users to scan them with their smartphones, bypassing email scanning tools entirely and connecting the user directly to a malicious site.

The United States: both target and source

In Q1 2025, the United States was the origin of 57 percent of global spam and the recipient of 75 percent of all malicious email traffic. This places U.S.-based organizations at the center of the global email threat landscape, both as victims and as unintentional participants in propagating spam and phishing.

Conclusion: defenses must evolve

The takeaway from both quarters of data is clear. Cybercriminals are shifting to low-signature, high-conviction attacks that bypass conventional email security by focusing on humans rather than systems. They are using AI, multilingual messaging, and cloud-based delivery to stay one step ahead.

Organizations must respond with equally intelligent defenses that focus not just on technical indicators, but on behavior, intent, and context. It is no longer enough to block malicious links or scan for known malware. Email security must evolve to recognize deception itself, using AI to detect and respond to the subtle tactics that define today's most successful attacks.

About the Author

Usman Choudhary is the Chief Product and Technology Officer at VIPRE Security Group. He leads product strategy, threat intelligence, and technology development for VIPRE's global cybersecurity and data protection portfolio.





Zero Trust: What Cybersecurity Experts Can Learn from Prisons

By Nick Kathmann, CISO, LogicGate

Just a few years ago, “Zero Trust” was the hottest buzzword in cybersecurity. In fact, it became so hot that every vendor wanted to use it—whether their solution adhered to Zero Trust principles or not. The result? The power of Zero Trust was significantly diluted. The term was manipulated by marketers and came to mean different things for different scenarios, creating confusion in the market and leaving customers uncertain about what Zero Trust actually means and whether it applied to them. Slapping the “Zero Trust” label on every security solution probably helped push product in the short term—but in the long term, it eroded confidence in Zero Trust principles themselves.

So, let’s clear it up with a simple analogy. Zero Trust is, basically, mirroring prison facility architecture. Zero Trust principles state that risk leaders should be designing their architectures with the goal of reducing the potential blast radius of a security incident. That requires tactics like micro segmentation, along with a strict and continuous approach to identity validation and data access privileges. High-security prisons are built on that same Zero Trust concept: access to the facility itself is extremely restricted, and even once inside, there are numerous security checkpoints, access barriers, and other safeguards designed to limit unauthorized movement or activity. Leaving one cell doesn’t immediately grant access

to all cells—and organizations need to take that lesson to heart if they want to limit risk within their digital infrastructure.

How Prisons Mirror the Zero Trust Approach

Access restriction and segmentation each play an important role in keeping prisons secure. Individuals cannot enter a prison facility unless they are on an approved visitors (or vendors) list. Those incarcerated within its walls cannot move between living units, the law library, gym facilities and other approved areas, unless granted permission or at specific, predetermined times. Correctional officers themselves require keys or keycards along with IDs and other verification and authentication methods to pass through security doors. For most institutions, access to the internet is highly restricted or prohibited, and all institutions are on high-alert to mitigate smuggled-in contraband. These are just a few examples, but the point is this: within the walls of a prison, movement—both physical and digital—is monitored, managed, and restricted.

When you break it down, the entire architecture of a prison facility is designed to protect the incarcerated, protect correctional staff, and protect the public. But incidents still happen. And when a disruption occurs, there are protocols and procedures in place to contain the situation, communicate the impact, and conduct a postmortem to ensure it doesn't happen again. To anyone in the cybersecurity field, that should sound pretty familiar—and it underscores the deep similarities between data security and traditional physical security. While it's easy to view the two fields as distinct, the truth is there is a lot that cybersecurity professionals can learn from their counterparts in the physical security field.

Applying Zero Trust Principles to Digital Environments

That basic lesson should help security and risk leaders think differently about how they build their network architecture. First, consider what Zero Trust actually means. Ideally, it means access is never granted by default—identities are not “trusted,” they need to continuously prove that they have the right to access certain systems and data. What's more, they should never have access to more data than they actually need and when they need it. This is referred to as the principle of least privilege: identities should have the minimum number of privileges needed to perform their essential functions, and nothing more. This helps significantly limit the impact of a potential breach: if a set of credentials is stolen, the attacker will only have access to a limited amount of data or systems, making it difficult for them to escalate the attack.

The parallel to a prison is clear. Incarcerated persons—and even guards—are not granted more access than they need. After all, if an incarcerated person could open every door in the prison with a single key dropped by a guard, that wouldn't exactly be ideal. In the real world, different doors would require different keys and different sets of credentials, and an incarcerated person attempting to access restricted areas would be repeatedly challenged to prove their identity—even if they somehow got ahold of a corrections officer uniform. There are multiple layers of defense, and none of them involve trust. If you can't prove who you are and why you should have access to an area, it simply won't be granted.

That is the important point—the one that underscores why Zero Trust is still relevant today. By segmenting different areas of the network and constantly challenging visitors to validate their credentials, organizations can effectively reduce the blast radius of an incident. This requires close collaboration between security teams and network/systems/identity architects, who can work together to analyze and quantify incident “blast zones” for potential business-impacting threats and apply threat modeling principles to determine trust boundaries. This will enable organizations to design better network, data, and access boundaries across different trust zones according to the potential impact of a security. By applying those principles across COTS, developed, and deployed architectures, they can make it harder for attackers to get in—and harder for data to get out.


Zero Trust Remains as Relevant as Ever

While risk is ever evolving, the foundational elements of security and mitigating threats are tried and true. Prisons aren't new—they have existed in some form for thousands of years. Admittedly, some need to modernize—but the security principles upon which they are built have not changed much over time. The importance of protection and emphasis on limiting surprises has never wavered, and even the most progressive and permissive prisons have multiple layers of security and authentication. Those are lessons that security teams and digital architects should take to heart. Embracing Zero Trust principles remains one of the most effective ways to limit risk in today's digital threat landscape.

About the Author

As CISO of LogicGate, Nicholas Kathmann brings a strong track record of delivering innovative security solutions for small to medium-sized businesses and Fortune 100 enterprises. He is an accomplished Chief Information Security Officer with over 20 years of extensive experience in IT and cybersecurity, specializing in application architecture and security management. Nick can be reached online at [LinkedIn](#) and at <https://www.logicgate.com/>





5 Steps to Move Beyond Vulnerability Discovery to Risk Remediation

By Chaz Spahn, Director of Product Management at Adaptiva

Vulnerabilities are growing rapidly in both volume and complexity. One [recent article](#) shows the National Vulnerability Database (NVD) continues to struggle to keep up with last year's unprecedented surge in reported vulnerabilities. Yet, despite cybersecurity teams detecting more issues than ever, breaches continue to slip through the cracks. Why? Because identifying vulnerabilities isn't a magic fix – it's only half the battle.

Threat actors are exploiting known vulnerabilities that organizations fail to remediate. According to the CSA's latest Annual Report on [Top Routinely Exploited Vulnerabilities](#), threat actors are most successful within the first two years after a vulnerability is publicly disclosed. After all, in cybersecurity, failing to promptly address problems is like having a state-of-the-art alarm system that never rings when it should.

Here's what to know to move beyond vulnerability discovery to risk remediation:

The Challenge Lies in Remediation

According to [research](#), 77% of organizations take over a week to deploy patches, while attackers typically exploit vulnerabilities within just five days.

Organizations know they should act fast, but the process is too manual, resources are insufficient, and the number of vulnerabilities seems insurmountable. Nevertheless, when teams hesitate and delay updates, they're left vulnerable to attacks.

Attackers only take days to cause harm, so closing the gap between vulnerability detection and remediation must be a top priority. It also means organizations need to expand their focus from detection to how quickly vulnerabilities are fixed.

Traditional Approaches to Vulnerability Management Hit Roadblocks

A common issue in remediation is that siloed IT and security teams traditionally cause bottlenecks in patching. While security teams identify threats, IT teams are responsible for maintaining operations and applying patches. As such, communication gaps increase, which slows down processes.

Many security teams still rely on static spreadsheets or CSV files to share vulnerability data with IT teams – formats that lack crucial, real-time context such as exploitability, asset exposure, and patch availability. As such, IT teams need to manually validate threats and search for remediation options, a time-consuming and error-prone process, particularly when facing thousands of vulnerabilities.

At the core of the issue is a fundamental disconnect: outdated, static data workflows are misaligned with the urgent need for real-time, actionable intelligence.

The Shift: 5 Steps to Prioritize Remediation

Unpatched vulnerabilities remain a priority for threat actors, and some known vulnerabilities have been exploited for years. This is largely because patching known vulnerabilities can be complex, time-consuming, and costly.

Timely patching, centralized patch management, and a sharp focus on time-to-remediation are essential pillars of a strong security posture. It's time for organizations to evolve their mindset—from simply counting vulnerabilities to asking the more urgent question: *"How fast are we fixing them?"*

Here are five steps for organizations to streamline vulnerability patching and remediation:

1. **Connecting processes – exposure management and vulnerability remediation:** A modern vulnerability management platform continuously assesses endpoint exposure, correlating common vulnerabilities and exposures (CVEs) and rating them by risk factors such as exploitability and business impact.

To drive effective remediation, these platforms should also integrate with patch management solutions that can automatically identify available patches and correlate them to specific vulnerabilities. This end-to-end visibility empowers administrators to detect, prioritize, and swiftly address exposures across the entire attack surface with greater precision and efficiency.

2. **Tailor remediation strategies to fit organizational needs:** Effective patching requires customization to overcome remediation challenges. According to [research](#), 64% of organizations cite coordination between detection and remediation as their biggest challenge when it comes to patching. As such, organizations should define patching strategies based on their unique needs, with humans setting the priorities and software executing deployments.
3. **Streamline patch deployments with automation:** Fast, accurate patching follows a controlled process. Automation helps coordinate approvals, testing, and update configurations before installation. Real-time vulnerability data ensures critical patches are deployed immediately, while other patches are applied based on severity.
4. **Stay proactive with adaptable controls:** Proactive remediation means staying ahead of patching issues. By monitoring user experience and leveraging controls, administrators can anticipate and resolve potential problems by providing notifications to varying teams about patch deployments and pausing, canceling, or rolling back patches as needed.
5. **Use compliance to your advantage:** Real-time reporting and monitoring provide visibility into the patching process, closing the loop from identification to remediation. Live progress updates help administrators track where updates are installed, monitor compliance improvements, and measure risk reduction.

Final Thoughts on Moving from Detection to Protection

Closing every security gap is a challenge, especially with the increasing number and complexity of vulnerabilities. To stay ahead of threats, organizations should focus on proactive remediation. This means fostering better collaboration between security and IT teams while leveraging automation to streamline the patching process.

With the right tools, organizations can continuously monitor their digital assets, ensuring vulnerabilities are addressed before they become exploits. And by incorporating autonomous patching, they can accelerate deployments, minimize disruptions, and ensure remediation keeps pace with detection – ultimately strengthening their overall security posture.

About the Author

Chaz Spahn is the director of product management at Adaptiva, a global leader in autonomous endpoint management, where he is tasked with evaluating customer feedback and conducting competitive analysis to identify and prioritize features that ensure our products remain best-in-class. He collaborates cross-functionally to keep all departments informed about product development and the status of new products and features. During his tenure at Adaptiva, Chaz has advanced from a Support role to developing an education program and platform for customers and partners. He then transitioned to the marketing team as a Product Evangelist and currently leads the Product Management team.





Agentic AI is Transforming Security – What Enterprise Leaders Need to Know

How Autonomous Systems Are Reshaping Threat Detection, Response, and Vulnerability Management

By Michiel Prins, Co-Founder & Senior Director, Product Management, HackerOne

Agentic AI is the latest evolution in artificial intelligence, bringing a new level of autonomy to digital systems. Unlike traditional AI, which relies heavily on user input, agentic AI operates independently, making decisions and taking actions based on real-time data. According to [Gartner](#), by 2028, one-third of enterprise software applications will incorporate agentic AI, up from less than 1% in 2024. This shift will enable 15% of day-to-day work decisions to be made autonomously, drastically changing industries, particularly in cybersecurity.

Agentic AI offers a new approach to threat detection, response, and vulnerability management. However, with its potential comes new risks. Enterprise leaders must understand both the benefits and challenges to make informed decisions about integrating agentic AI into their security strategy.

Understanding Agentic AI vs AI Agents

At its core, agentic AI consists of autonomous agents capable of reasoning, learning, and taking independent action to achieve specific goals. Unlike traditional AI models that generate responses based on prompts, agentic AI engages in iterative workflows, adapting to dynamic environments in real time.

A key distinction exists between agentic AI and AI agents. While AI agents perform individual tasks, agentic AI represents a broader system of interacting agents capable of sophisticated decision-making and self-improvement.

This shift marks a departure from conventional automation. Traditional AI solutions often rely on predefined rules and models, while agentic AI can assess complex situations, weigh multiple factors, and act with minimal human intervention. In cybersecurity, this capability is invaluable for responding to evolving threats and helping with remediation.

The Benefits of Agentic AI in Cybersecurity

Enhanced Threat Detection and Response

Agentic AI accelerates cybersecurity operations by reducing the Mean Time to Detect (MTTD) and Mean Time to Conclusion (MTTC). Through real-time anomaly detection and automated incident response, agentic AI can swiftly identify and neutralize threats before they escalate. Lowering MTTD means detecting threats faster, minimizing the time attackers have to exploit vulnerabilities, while reducing MTTC ensures threats are mitigated swiftly before they cause major disruptions.

Moreover, agentic AI minimizes human error by automating repetitive tasks, allowing security professionals to focus on complex investigations. By eliminating alert fatigue and prioritizing critical threats, security teams can operate more effectively.

Scalability and Cost Efficiency

As cyber threats grow in volume and complexity, hiring additional security analysts becomes increasingly difficult and expensive. Agentic AI provides a scalable solution by automating key security functions, reducing the need for large Security Operations Center (SOC) teams while maintaining high levels of protection.

How Security Will Evolve with Agentic AI

Traditional security solutions rely on human-led decision-making, often reacting to threats after they occur. In contrast, agentic AI-driven security will shift from reactive to proactive threat management.

For example, current AI-driven security tools, such as Large Language Models (LLMs), primarily assist in generating responses or analyzing data. With agentic AI, security tools will move toward self-sufficient operations, autonomously managing cyber risks and adapting to new attack vectors. This shift will redefine security as a service, transitioning from software as a service (SaaS) to service as a software.

Applying Agentic AI in Security

Penetration Testing

Agentic AI can significantly enhance penetration testing (pentesting) by automating the initial stages of vulnerability discovery. AI agents can rapidly scan attack surfaces, and identify common security flaws. This enables human pentesters to focus on more sophisticated attack vectors, leading to deeper and more comprehensive security assessments.

Threat Detection and Response

Security Operations Center (SOC) teams are often overwhelmed with alerts, many of which turn out to be false positives. Agentic AI can autonomously investigate these alerts, distinguishing genuine threats from benign anomalies. By handling routine investigations, agentic AI allows human analysts to concentrate on high-priority incidents.

Additionally, AI agents can be scaled instantly, effectively doubling a security team's capacity without hiring additional staff. This automation enhances efficiency while reducing operational costs.

Vulnerability Management

Prioritizing vulnerabilities has long been a challenge for security teams, requiring extensive context gathering and analysis. Agentic AI can automate this process by reviewing historical vulnerability data, assessing potential risks, and recommending remediation strategies. This streamlines the workflow, ensuring that security teams focus on the most critical threats.

Balancing the Risks and Rewards of Agentic AI

While agentic AI offers significant advantages, it also presents new risks. Organizations must implement safeguards to prevent unintended consequences and security vulnerabilities.

Guardrails and Human Oversight

Excessive autonomy can be problematic. AI agents must be designed with oversight mechanisms to prevent misjudgments or unintended actions. Human-in-the-loop workflows ensure that critical decisions receive human review before execution.

A recent [legal ruling involving Air Canada's AI chatbot](#) highlighted the importance of accountability. Companies cannot dismiss AI-generated errors as mere software glitches. Instead, they must assume responsibility for their AI systems and establish clear governance policies.

Mitigating Prompt Injection Attacks

Agentic AI's ability to gather and process vast amounts of data also introduces security risks. Malicious actors can manipulate AI agents through prompt injection attacks, steering them toward unintended actions. This threat underscores the need for robust security protocols and continuous monitoring.

Data Access and Privacy Concerns

Organizations must carefully control AI access to sensitive data. Without proper restrictions, AI agents could inadvertently expose confidential information. Security teams should implement strict data governance policies to mitigate these risks.

Preventing Jailbreaks

AI safety mechanisms are essential to prevent unauthorized access or abuse. Attackers can exploit vulnerabilities in AI models to bypass restrictions, leading to unintended actions. Implementing strong security frameworks and regular audits can help safeguard AI systems from exploitation.

The Future of Agentic AI in Cybersecurity

The integration of agentic AI into cybersecurity is inevitable. As AI capabilities continue to progress, businesses must embrace these advancements while implementing risk management strategies.

According to [Grand View Research](#), AI adoption is expected to grow at an annual rate of 36.6% between 2023 and 2030. By 2028, Gartner predicts that 15% of day-to-day enterprise decisions will be made autonomously through agentic AI.

For enterprise leaders, the key takeaway is clear: adopting agentic AI responsibly will be crucial for maintaining a competitive edge in cybersecurity. By balancing innovation with security best practices, organizations can harness the full potential of agentic AI while mitigating its risks.


Conclusion

Agentic AI represents the next frontier in AI and cybersecurity, offering enhanced threat detection, operational efficiency, and scalability. However, responsible implementation is critical to avoid unintended security risks. By incorporating human oversight, securing AI workflows, and prioritizing ethical considerations, organizations can confidently integrate agentic AI into their security strategies.

As agentic AI reshapes the cybersecurity landscape, enterprise leaders must stay informed and proactive. The future of cybersecurity lies in leveraging AI-driven solutions while maintaining a strong commitment to security and ethical governance.

About the Author

Michiel Prins is a Co-Founder and Senior Director of Product Management at HackerOne, the cybersecurity company dedicated to eliminating vulnerabilities through continuous testing. He is an information security expert, researcher, hacker, and developer. Michiel has been finding critical software vulnerabilities in technology for over 10 years. Prior to founding HackerOne, Michiel co-founded a successful penetration testing company that worked on projects for trusted organizations from government institutions to top technology companies, including Twitter, Facebook, Evernote, and Airbnb, among others. Michiel regularly presents on vulnerability disclosure and security research projects regarding security management, privacy, and web application infrastructure. Michiel graduated with a B.S. in Computer Science from Hanze University Groningen.



Michiel can be reached online on [LinkedIn](#) and at our company website <https://www.hackerone.com/>



AI is the New Attack Surface

What CISOs Must Know

By Patrick M. Hayes, Author & Security Strategist, Integrated Assurance LLC

AI is no longer a technical experiment operating on the fringes of business operations. AI is now part of core business workflows, decision-making processes, customer experiences, and major strategic efforts across almost every industry. Its strength lies in how it learns, adapts, and scales with ease. But those same strengths also create new risks that most conventional security and governance approaches simply aren't built to handle.

As organizations rush to capitalize on advantages of AI, they typically underestimate the breadth and depth of the threat surface it creates. AI may be a powerful tool for innovation, but it's also a new vector for exploitation, failure, and regulatory exposure. The systems we are integrating into our businesses can behave in ways that are difficult to predict, harder to control, and nearly impossible to explain once they go wrong.

The time has come for enterprise security leaders to reposition themselves as enablers of trust at scale and not just risk mitigators. This shift demands a new approach. It requires Integrated Assurance, a model that treats security as a continuous and embedded discipline, aligned with innovation, not opposed to it. This article outlines the key dimensions of AI as a threat vector and presents a practical path forward for enterprises committed to building responsibly in an AI-powered world.

AI Risks Are Unlike Traditional Threats

Traditional security frameworks were designed for deterministic systems with fixed logic and understood failure modes. AI systems, especially those based on deep learning, behave differently. They evolve over time, generate outputs based on statistical associations, and often lack clear reasoning behind their decisions. This makes the output of AI systems difficult to validate, monitor, and govern using legacy methods.

For example, when AI systems are trained on biased data, they can make discriminatory decisions even if the code itself seems fine. Generative models might create content that looks accurate but is actually false. Once these systems are built into everyday operations, they can quietly disrupt performance, violate ethical standards, or create legal risks across the organization.

Compounding the problem is speed. AI adoption is outpacing governance maturity in most organizations. Business units experiment with third-party tools, developers prototype with unvetted models, and leadership teams greenlight AI initiatives with limited visibility into risk implications. Security and assurance teams are often brought in after deployment, too late to influence design or validate safeguards. This results in a fragmented ecosystem where risk accumulates quietly and explosively.

Prompt Injection and Semantic Exploits

Unlike traditional software, AI models interpret natural language. This creates an entirely new form of prompt injection attack. By embedding malicious or manipulative instructions into user-facing inputs, attackers can bypass policies or trigger unintended behavior in AI systems.

Consider a scenario where a generative assistant is integrated with enterprise workflows. If an attacker enters a prompt that mimics internal authority or overrides prior instructions, the model may approve transactions, escalate tickets, or disclose sensitive information. These systems interpret meaning, not syntax, which makes conventional input validation ineffective.

Prompt injection exploits are difficult to detect because they look like normal conversations. They blend into the flow of communication, relying on the model's willingness to comply with what appears to be a legitimate request. Security teams must begin thinking about inputs not just as data, but as potential command surfaces. Defending against this class of attacks requires runtime controls, model constraints, user education, and forensic logging at the prompt level.

Synthetic Media and the Collapse of Trust

AI generated content, like deepfakes, voice clones, and fake documents, is forcing us to rethink identity, communication, and trust. It's now easier than ever to impersonate leaders, spread false information, and manipulate narratives. The result is a growing confusion between what's real and what's fabricated, not just within organizations but across public conversations as well.

What makes synthetic media dangerous is not just its realism. It is the ease and scale with which it can be produced and distributed. A deepfake voice call or video message can now be generated in minutes and deployed in social engineering campaigns that bypass technical defenses. The result is a fundamental erosion of confidence in audio, visual, and textual communication.

Enterprise security must now account for the authenticity of messages, not just their source. This means putting media verification tools in place, setting up separate channels for secure authentication, and having a clear crisis response plan ready in case of synthetic impersonation. Training also needs to go further, teaching employees and stakeholders how to spot fake or manipulated content before it can do real damage.

Shadow AI

Across organizations, employees are using AI tools outside of sanctioned environments. They paste proprietary content into generative interfaces, rely on code written by AI without validation, and use free third-party models to make business decisions. This shadow AI phenomenon is growing rapidly and silently.

The risks are significant. Sensitive data may be stored or reused by external providers. AI generated outputs may contain hallucinated content or introduce security vulnerabilities. Decision logic may vary across teams, leading to inconsistent and biased outcomes.

The answer isn't to ban AI completely. Instead, organizations need to build safe environments where innovation can happen with the right oversight and controls. That starts with using approved AI platforms, setting clear policies based on real business needs, and closely monitoring how AI is being used to catch any issues early. Security teams should see their role not just as gatekeepers, but as partners in helping the organization use AI in a smart, responsible way.

Model Inversion and Extraction

AI models are no longer just tools running quietly in the background. They can now be poked, prodded, and manipulated through the same public interfaces meant for users. Attackers are finding ways to reverse-engineer these models, pull insights about their training data, or even copy their core functionality outright.

The consequences are serious. These types of attacks can lead to data breaches, stolen intellectual property, or a noticeable drop in model performance. What makes it worse is that traditional security tools often miss them. From the outside, the model still appears to work normally, so no alarms are triggered.

To stay secure, organizations need to treat AI systems as core components of their security environment. This involves putting clear safeguards in place, like managing who has access, limiting how often the systems can be used, applying privacy-focused techniques, and closely monitoring for any suspicious user activity. It's also important to go deeper than just the application layer. Security checks should include the AI model itself, since it can be a potential target for attacks.

Regulation Is Accelerating

Governments around the world are moving quickly to put rules in place for AI. In the EU, the AI Act is taking shape. In the US, executive orders and state laws are starting to set clear expectations. Across these efforts, the message is the same, *“AI needs to be transparent, accountable, and guided by human judgment”*. Businesses that ignore this shift are taking a serious risk.

Meeting regulatory demands today takes more than checking boxes. It means keeping an up-to-date inventory of AI systems, clearly documenting where data comes from and how decisions are made, and keeping a close eye on how models behave over time. If a company relies on third-party AI tools, contracts need to spell out the right to audit and set clear performance expectations.

Rather than seeing regulation as a burden, security leaders should view it as a chance to improve governance. Taking a proactive approach to emerging standards can build trust, set the company apart from competitors, and ensure its prepared for whatever the future brings.

The Strategic Response

Integrated Assurance provides a framework for governing AI across the full enterprise lifecycle. It's an operational model that aligns engineering, risk, compliance, legal, and security around a shared mission of building trustworthy AI.

This model begins with visibility. Every AI system must be registered, classified, and understood. Next comes consistency. Policies, templates, and evaluation criteria must be standardized so that assurance can scale across diverse use cases. Then comes adaptability. Governance mechanisms must respond to changes in model behavior, business needs, and external regulations.

Integrated Assurance also emphasizes culture. Trust is not created through documentation alone. It is created through behavior from leaders modeling transparency, teams collaborating across functions, and every stakeholder taking ownership of risk.

Security leaders have a unique opportunity to lead this shift. They understand the complexity of modern systems, navigate nonstop regulatory pressure, and stay alert to how fast threats can change. When they adopt Integrated Assurance, they stop being seen as the ones who hold things back and start being valued as the people who make trust possible.

A Call to Leadership

The future of organizational resilience won't be shaped by how many AI tools a company rolls out. What will truly matter is how well those systems are governed. The organizations that succeed in this next phase won't necessarily be the fastest or the most cost efficient, they'll be the ones people trust.

Security has to lead from the front. That means moving past silos, aligning closely with business goals, and making sure assurance is built into every part of the AI lifecycle, from design to deployment. It starts

with a shift in mindset and anticipating failure before it happens, instead of reacting once things go wrong. It also means building systems that aren't just advanced or efficient, but that people can understand, audit, and trust to reflect the organization's values.

AI is rewriting the risk landscape, and with it, the responsibilities of the security architect. This role is no longer just about safeguarding data or locking down infrastructure. It's about protecting the quality of decisions, the transparency of communication, and the trust that users, employees, and stakeholders place in the entire system.

That's where Integrated Assurance comes in. It gives us a way to keep AI within the guardrails of governance. And more importantly, it gives us a path to build a future that moves fast and stays grounded in trust.

About the Author

Patrick M. Hayes is the Author of Integrated Assurance: Unified Risk Strategy. He is a certified enterprise security architect with over 25 years of experience at the intersection of cybersecurity, IT operations, and business transformation. Patrick is a recognized leader in aligning enterprise security with emerging technologies, specifically artificial intelligence. His work focuses on helping global organizations modernize governance, risk, and assurance strategies to meet the challenges of AI-driven automation, synthetic threats, and regulatory change. Patrick speaks internationally on AI as a threat vector, trust-centric architectures, and operational resilience in an era of intelligent risk.



Patrick can be reached online at www.integratedassurance.co or connect with him on LinkedIn at <https://www.linkedin.com/in/phbalance/>.



Artificial Intelligence Has Changed What Cyber Defense Means for the Retail Industry

By Jason Lewkowicz, Managing Director – Cybersecurity, America's Retail Sector Lead, Ernst & Young LLP

To secure their organizations from an evolving artificial intelligence (AI) cyber threat landscape, retail CISOs must prepare for external attacks and establish consistent internal safeguards.

It's not clear whether retail industry chief information security officers (CISOs) are ready for the reality of how AI has changed the cybersecurity landscape for their organizations. In 2024, it was estimated that the retail sector is targeted by over 500,000 AI-driven cyber attacks each day. Despite the seeming nonstop wave of threats, in a recent EY survey, only 40% of retail CISOs perceive AI as a potential security threat to their organizations. And even those leaders may not be viewing the generative AI (GenAI) cyber defense threat through the right lens.

Beyond the headlines

Media reports about AI and cybersecurity usually focus on how cyber criminals are using emerging AI tools to enhance attacks on companies, so we can be confident retail CISOs understand that AI has

revolutionized the speed and complexity of digital information reconnaissance and deepfake-enabled smish, vish and phish attacks. Attackers are using AI to create deepfakes, mirroring voices and writing styles to deceive help desks, support systems and employees elsewhere in the enterprise. Common attacks include posing as a senior leader and ordering a junior employee to purchase gift cards to help close a deal. With AI's ability to scrape vast amounts of information about not only CISOs and other key personnel in the enterprise but their families and acquaintances, the success rate of these attacks has dramatically increased. Historically, digging up enough information to fuel an attack would require a team of hackers; now, AI agents can achieve this autonomously and efficiently.

The standard narrative of hackers breaching security through AI tools only presents a part of the picture, which prevents CISOs from understanding the true, multifaceted ways that AI could contribute to breaches. To effectively mitigate risks, CISOs must adopt a holistic perspective and see not just what the bad guys are doing with AI but what the good guys in retail organizations are doing and should be doing to integrate AI into their systems and use it to strengthen cybersecurity.

Clarifying the nature of the threat

You don't need a hacker or an external attack for AI to pose a cybersecurity threat to a retail organization. When CISOs view GenAI only as a tool for extracting sensitive data, they overlook the possibility that AI can also inadvertently *expose* sensitive data, which can do as much harm to operations as an attack. Misconfigured AI systems can leak proprietary information, affecting consumer trust and competitive advantage. For instance, without proper restrictions, consumers could ask a retailer's propriety chatbots to compare the company's product line with competing products, turning a potential sale into a win for the competition.

In addition to harming businesses through offering insights into competitor products, AI chatbots set up with improper security guidelines could access and reveal confidential information about an organization or its personnel. Imagine a scenario where an organization-wide AI system connects previously disparate data sources, such as human resources data about salaries and so forth, across the entire organization, including a public-facing chatbot. All it takes is one curious online shopper writing a simple prompt for employee X's compensation to go public.

Harnessing AI for cyber defense

Retail organizations must think defensively about deploying AI. If you're not already on this AI journey, it's essential to start, or risk being left behind. The scenarios described above underscore the necessity of robust AI deployment strategies. The silver lining to the AI threat to cybersecurity is that while AI can make data more vulnerable, it also has great potential to strengthen cybersecurity. By leveraging AI for threat detection, anomaly monitoring and streamlined reporting, retail organizations can enhance their cyber defense strategies. AI can automate routine tasks, optimize threat hunting and facilitate efficient security operations.

But the security benefits of GenAI can only be realized when organizations build a solid foundation for how the technology should work throughout the enterprise. Before deploying AI, organizations must establish robust policies for AI deployment and standards for AI use. Employees should have full clarity on acceptable use of AI tools and technologies, including knowing which behaviors are restricted and what disciplinary actions would be for violating those restrictions. These policies not only define operational boundaries but also facilitate corrective measures in case of violations.

Defining models and cost management

Deciding which AI models to deploy is pivotal to establishing a firm foundation for security. Organizations aligned with tech giants may have different preferences, influenced by cost and compatibility. Security often functions as a cost center, requiring strategic decisions on model usage to avoid financial pitfalls. Clarity on model selection facilitates effective cost management and consistent outcomes.

As you roll out these language models, you must train them on what information they can consume from your enterprise and what they cannot. Depending on the role of the agent, what information do you want it to give to the requester?

AI systems must be continuously validated to verify secure operations just as security would validate that controls and stop-gap measures work correctly in non-AI systems. Organizations must establish checks and balances, confirming that AI systems create secure code and produce secure outcomes. This diligence prevents unintended data exposure and reinforces cybersecurity frameworks. Organizations must ensure that third parties can't access customer information that store associates or online stores use to personalize shopping experiences, as that information is a competitive advantage that could help other businesses.

Conclusion

Through informed decision-making and strategic planning, retail CISOs can turn potential threats into opportunities, transforming GenAI from a perceived risk into a valuable asset. Comprehensive policies and rigorous validation processes are essential to enhance cybersecurity. Consulting with EY professionals can provide valuable guidance in securing AI deployments and protecting against evolving threats. By leveraging AI responsibly, retail CISOs can transform perceived threats into opportunities, strengthening their defenses and safeguarding their businesses.

The views reflected in this article are the views of the author and do not necessarily reflect the views of Ernst & Young LLP or other members of the global EY organization.

About the Author

Mr. Lewkowicz is part of EY's Technology Cyber Consulting Service Line focused on the Retail vertical. With more than 25 years of experience managing, leading and transforming complex global programs, he works to partner with EY's clients by advising, accelerating and running their programs.



Prior to his role at EY, Lewkowicz served as Executive Vice President and Chief Services Officer for Optiv. He was accountable for a global organization of more than 1,600 employees serving a wide variety of client challenges. He also was responsible for Optiv's India business, Partners & Alliances, and was the Client Advisory Board chair.

Prior to Optiv, Lewkowicz was the chief information security officer (CISO) for Cognizant Technology Services supporting their global business, including more than 50 acquisitions. Lewkowicz joined in 2020 as the pandemic was ramping up and was instrumental in navigating the firm through its Maze breach which included a full rebuild and recovery of the enterprise IT systems.

Prior to his role at Cognizant, Lewkowicz was a deputy CISO at Accenture where he managed a global team of more than 100 people operating across seven countries. His primary responsibilities included cyber response, forensic investigations including e-discovery, data loss prevention (DLP), threat hunting, red team exercises, cyber threat intelligence and cyber metrics and automation. Lewkowicz developed Accenture's cyber response program and all functions within it over a 17-year period. During his tenure, he was also instrumental in the creation of Accenture's overall information security program.

Lewkowicz holds certifications from ISACA, ISC2 and Open Text/Guidance Software. He has presented at industry-recognized conferences and local law enforcement information sessions over the past 10 years. He has served on advisory boards for Symantec, McAfee, Seculert, Digital Shadows, FireEye/Mandiant, ArmorBlox and Immersive Labs. He received his undergraduate degree from DePaul University in Chicago.



Branding Beyond the Breach: How Cybersecurity Companies Can Lead with Trust, Not Fear

By Lucia Barbato, CEO, Ilex Content Strategies

Cybersecurity breaches can instantly derail operations, damage reputations and reduce customer trust and interest. As cyberthreats grow in sophistication and complexity, cybersecurity firms need strategic messaging and positioning that leads with confidence, not fear, to stand out, build trust and grow in new markets.

Next generation technologies such as cloud, 5G and AI applications are rapidly growing in use, scale and complexity worldwide. To support their adoption, B2B companies need cybersecurity services that they can trust to support their business operations seamlessly. According to the [UK Government](#), 41% of businesses sought cybersecurity guidance from out-of-house sources in 2024, directly indicating a lack of in-house expertise and the requirement for strong cybersecurity approaches across a wide range of sectors.

As cybersecurity threats continue to evolve, cybersecurity companies must shift their messaging and tone away from fear-based strategies to stand out in a crowded market. They need messaging that highlights trust, resilience, compliance and creates positive real-world outcomes, rather than fearmongering based on past experiences or threats.

The Pitfalls of Fear-Based Messaging

Cyberthreats are growing in intricacy day by day. AI-powered attacks such as deepfakes help to bypass security measures, and increased connectivity such as cloud and the Internet of Things (IoT) open more vulnerabilities. On top of this, the industry has seen ransomware shifting towards triple extortion tactics, and nation-state actors are increasingly targeting government agencies, defence and critical infrastructure.

The almost constant stream of cyberattack headlines in the news only highlights the importance for cybersecurity companies to ensure their messaging is creating trust and confidence for B2B businesses. With the global cybersecurity market volume expected to reach \$262.29bn by 2030, according to [Statista](#), the challenge for cybersecurity companies lies in being able to differentiate themselves against competitors - it's not just about telling B2B companies to trust them, but showing them why.

It is easy to take issues such as AI- powered attacks and triple extortion tactics and create fear-based messaging in hopes of capturing attention. However, when cybersecurity companies endlessly recycle breach risks as reasons to do business, it can overload prospective clients with the dangers and cause them to disengage. It also minimises cybersecurity services down to being solely reactive, rather than proactive and preventative. By following fear-based messaging, cybersecurity companies are blending in, *not* standing out.

C-suite decision-makers value cybersecurity that will ensure the safe and seamless completion of mission-critical operations. These include features such as strong uptime, high redundancy and low latency. Cybersecurity firms need to move away from using technical jargon and threat factors in their messaging, which don't add value to B2B business propositions. Instead, highlighting factors such as business continuity and risk mitigation will position cybersecurity firms as industry leaders, ensuring B2B organisations feel seen and heard regarding their cybersecurity risks and fears.

Standing Out in a Crowded Landscape

To break away from fear narratives, cybersecurity companies need to align their messaging with what B2B businesses really want: strategic guidance, business continuity and actionable, measurable value.

Cybersecurity messaging should reflect this - but how?

- **Lead with Trust and Confidence, Not Fear**

Cybersecurity is not just a risk minimiser, but an asset and enabler. Strong cybersecurity ensures uptime, regulatory compliance and international scalability and this should be championed in messaging. Proving service quality and reliability over pre-existing threats positions cybersecurity companies as proactive and preventative in a landscape full of cyber risks, which ultimately keeps business operations secure and seamless for their clients.

- **Reduce Fear with Facts and Outcomes**

Transparency builds trust and credibility. Cybersecurity companies should be open about the true reality of risks and highlight proactive elements, such as how their technology combats threats and why this will support business operations. Most importantly, the focus should be on outcomes, not causes or technical terms. Messaging should focus on how attacks will be minimised, *and what this means for businesses*.

- **Keeping Up with Global Demands**

In today's globalised, multi-sector economy, cybersecurity messaging isn't one-size-fits-all. Messaging should vary based on factors such as geographic, industry and cultural differences. For example, healthcare verticals should prioritise patient data protection, whilst financial verticals may value fraud prevention and protection more. By mirroring the outlooks of global B2B companies in this way, confidence and trust is built in cybersecurity providers. This in turn will successfully secure business for cybersecurity firms competing in a saturated market.

- **Be a Strategic Advisor, Not Just a Vendor**

To navigate the complexities of cybersecurity, B2B businesses need a partner to guide them, not just sell to them. By including thought-leadership, education initiatives, consultation services, partnerships and customised strategies into a cybersecurity company's messaging and offering, it highlights their authenticity, credibility and reliability. Showing prospective clients that cybersecurity companies are experts in the industry and can offer operational stability and long-term scalability through their expertise will secure business quickly and effectively.

The Future of Cybersecurity Messaging


The cybersecurity landscape is wide and complex, and the market will only continue to diversify as threats evolve. Cybersecurity organisations need messaging that shows they can support businesses to expand in new sectors, communicate complex offerings clearly and become the optimal solution for risk-conscious enterprises.

A narrative that is clear, confident and customer-oriented rather than risk-focused will translate the complexity of cybersecurity into valuable takeaways. It's branding beyond the breach that matters for the success of cybersecurity businesses in today's rapidly evolving landscape.

About the Author

Lucia Barbato is the CEO of Ilex Content Strategies, a global B2B marketing agency working within the tech and telecommunications sectors. Lucia is an expert in content marketing, strategy, and PR. Lucia can be reached online at info@ilexcontent.com and on LinkedIn at [Ilex Content Strategies](#), and at our company website <https://www.ilexcontent.com>





Building Responsible AI: Staying Ahead in a Regulated World

By Paul Davis, Field CISO, JFrog

The entire world of technology is abuzz about AI/ML. It's arguably the most disruptive technology to society since the smartphone. Enterprises recognize they must adapt; Gartner estimates that by 2027, over 90% of new software applications will include machine learning models. While this rapid advance is fueling quantum leaps in innovation, it also ignites increasing scrutiny from regulatory bodies worldwide, demanding unprecedented levels of rigor, transparency, and accountability from development teams deploying AI/ML in production environments.


We'll discuss how governments are responding to put guardrails around AI/ML development, and how you can maintain a proactive stance to protect your business from regulatory fallout.



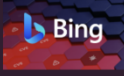







How Regulators View AI/ML

AI/ML development happens at an unprecedented breakneck speed, and inherently requires a vast volume of large datasets, complex algorithms, and continuous model training, significantly widening the scope of risk assessment in software development. Given the sheer power of AI/ML, it's no surprise regulators are intensifying their scrutiny.

Beyond traditional concerns like code quality and data security, regulators are now assessing intelligent agents for their potential harmful impacts, ethical implications, and societal consequences. For example, regulators have data privacy and intellectual property concerns around how AI/ML models are trained, and are worried about the dangers they create, like discriminatory outcomes (algorithmic bias) or misinformation (AI hallucinations).

AI/ML Regulations are Coming, Now



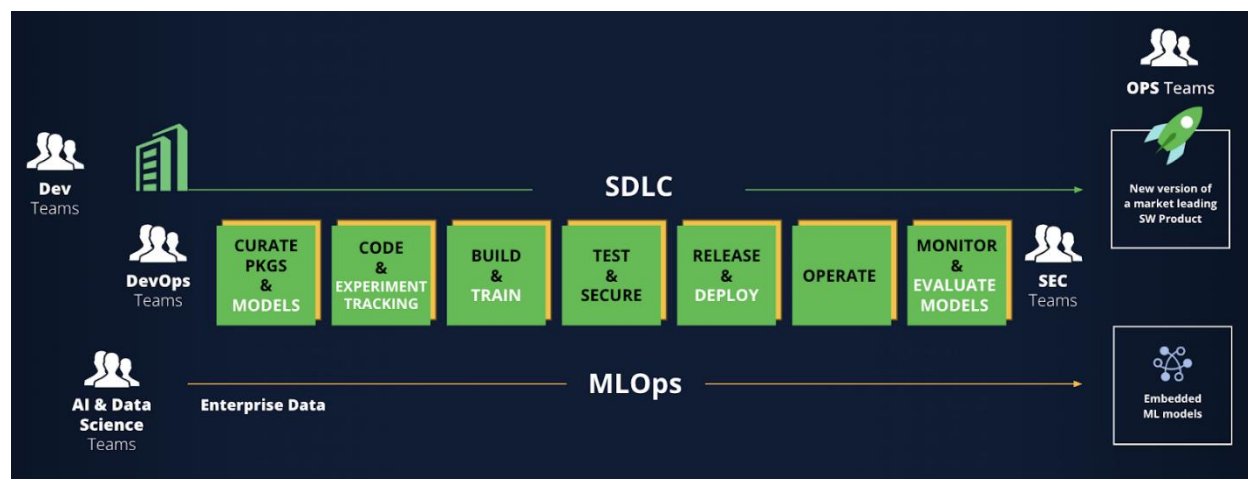
	EU AI Act Concerns 2023	
	Bing Chat Vulnerability 2023	
	ShellTorch 2023	
	DPD Chatbot Exploited 2024	
	Malicious Hugging Face Models 2023-2025	

Over in Europe, regulators have formalized comprehensive regulations around the need to thoroughly assess and certify your different AI-based solutions before deployment. There are serious ramifications associated with failing to comply with the European Union Artificial Intelligence Act, which classifies AI systems based on risk levels and imposes strict requirements accordingly. There are significant financial penalties, which can be extremely substantial for enterprise organizations, potentially reaching up to 6% of global annual revenue for the most severe violations - figures that could translate to billions of dollars for larger companies.

In the US, AI regulations are being defined at the state level (California, Colorado, New York, Texas, and Washington) and federal level (Executive Order 14110).

Other countries are also suitably concerned and introducing their own requirements, building around these existing laws as well as ISO 42001. With all these demands to demonstrate compliance, the potential headache associated with a whole new supply chain model can be daunting.

The Solution: An Integrated Approach to Machine Learning Operations (MLOps)



Imagine a world where the system would not let you promote a software release into production unless it had passed **all the tests** (yes, I mean all the tests!). Exceptions could be recorded, tracked, and the approvals documented in a single system.

Continuous Compliance Automation offers the solution. By applying controls and enforcing regulations from beginning to end—from initial design to production release—we can eliminate the need for point-in-time checks (i.e., audits). We integrate [evidence](#) from all your tools and procedures into a single data source of truth, enabling everyone (developers, AppSec teams, security personnel, auditors, and business owners) to see how the software complies with their specific regulatory requirements. This is also a great benefit to ML engineers and data scientists, who can have peace of mind knowing that they're working with trusted models that have been vetted and conform with policy.

Implementing an approach that combines consistent tools and processes with a reliable path to production creates a trusted environment that automatically generates information demonstrating adherence to regulations and compliance obligations. As some of my fellow CISOs are beginning to recognize, the most effective way to integrate security is to automate it through a methodical process.

Leveraging JFrog for Trusted AI/ML Development

AI, ML, LLM, and GenAI are here to stay. With the upcoming challenge of agentic AI, we need to build the foundations of a secure approach to managing risk. This includes addressing traditional concerns like vulnerabilities, personally identifiable information (PII), and business risks, while also developing the flexibility to adjust to new demands as the world confronts emerging threats generated by this new intelligence landscape. Start today, since tomorrow won't wait.

In the JFrog Platform, we have a full solution for [AI/ML model lifecycle management](#), which provides a trusted environment for AI/ML developers to build and fine-tune models. With JFrog ML, you gain governance over your ML development, with end-to-end visibility and control into AI model deployment,

including usage and permissions. You can also leverage [Evidence Collection](#) to collect evidence of every step taken to advance AI models into production.

Need help getting started? [Take a tour](#), or [book a demo](#) with the MLSecOps experts at JFrog!

About the Author

Paul Davis is an experienced IT Security Executive who, as Field CISO at JFrog, works to help CISOs, IT execs and security teams, enhance protection of their software supply chain. Additionally, he advises IT security startups, mentors security leaders, and provides guidance on various IT security trends. Paul also spends his time exploring the latest technologies, DJing, reading, and boating. <https://jfrog.com/>





Wiper Malware: What Federal, State and Local Agencies Must Know to Protect Mission-critical Systems

By Dirk Schrader, VP of Security Research and Field CISO EMEA, Netwrix

When malicious code is engineered not for ransom but for ruin, the very continuity of government is at stake. Wiper malware has emerged as a favored weapon of nation-state actors and hackers intent on erasing data, disrupting essential services, and undermining public trust. From the NotPetya outbreak that crippled ministries in Ukraine to the Shamoon campaign that paralyzed regional energy agencies, recent incidents show how quickly a single destructive payload can ripple through the networks that power citizen services, defense operations, and critical infrastructure. Understanding how wipers work – and how to stop them before they detonate – is now a core mandate for every agency CISO and IT security leader.

Understanding the Nature of Wiper Malware

Wiper malware is increasingly being used in cyberwarfare, hacktivism and politically motivated attacks. One of the most vivid examples is the [NotPetya attack in 2017](#), which impacted businesses and government agencies worldwide and caused an estimated \$10 billion in total damages. Interestingly,

NotPetya disguised itself as a ransomware attack, leaving a ransom note on the screens of infected devices, although decryption was impossible. This approach intends to distract incident responders.

Attacks like NotPetya highlight the immense financial and operational toll that wiper malware could impose nearly a decade ago, and the risk to organizations has only grown in the intervening years, with known examples like Whispergate (2022) and Apostle (2012-2025).

Wiper attacks aim to irreversibly erase data, disable critical infrastructure and cause chaos, rather than to demand a ransom or steal information. This makes them a preferred tool for adversaries such as hackers and nation-state actors engaged in cyberwarfare, ideological conflicts or political retaliation. Unlike financially motivated cybercriminals, these attackers prioritize causing disruption, damaging reputations or advancing strategic national interests.

Organizations most vulnerable to wiper malware attacks include government agencies, defense contractors, critical infrastructure operators and enterprises in politically sensitive regions. Those at the highest risk are entities that are caught in geopolitical conflicts or that store and process sensitive data that adversaries want to eliminate rather than monetize.

Real-World Impacts of Wiper Malware

A well-known example of wiper malware wreaking havoc is the 2014 [Shamoon](#) attack on Sony Pictures. [This attack](#), allegedly orchestrated by North Korea, targeted Sony in retaliation for its satirical movie *The Interview*. The Shamoon virus was designed to erase data on infected machines, making recovery nearly impossible. Beyond the direct data loss, the attack led to the leak of Sony's sensitive internal communications, unreleased films and employee data. Estimates [suggest](#) Sony incurred losses of at least \$15 million from this cyberattack, not to mention the long-term reputational damage.

More recently, wiper malware has become a key component of cyberwarfare. In international conflicts, state-sponsored attackers use destructive malware to disrupt critical services, cripple financial institutions and weaken national infrastructures. Such attacks can have far-reaching consequences, including economic destabilization and public safety risks.

8 Best Practices for Facing Wiper Malware Attacks with Confidence

Defending from wiper malware attacks requires a robust cybersecurity strategy similar to protecting against other sophisticated threats. The core principles remain the same: preventing initial intrusion, ensuring rapid response to minimize the impact of successful breaches, and speedy recovery to mitigate operational disruption. Here are the most effective practices:

- **User awareness and training** — Educate employees on how to recognize phishing attempts, social engineering tactics and other techniques attackers exploit to enter the network in order to deploy wiper malware.
- **Regular software updates and patch management** — Unpatched vulnerabilities are a common entry point for malware. Ensure all systems, applications and network devices are promptly

updated with the latest security patches. Where feasible, allow automated patching for security updates, like in Windows environments.

- **Access control based on the least privilege principle** — Restrict user access based on job roles and limit the privileges granted to service accounts. Ideally, grant elevated access only when necessary for particular tasks. Minimizing permissions dramatically reduces the reach of both malicious insiders and adversaries who compromise accounts.
- **Network segmentation** — Implement network segmentation to contain infections and prevent malware from spreading across critical systems. In particular, be sure to isolate environments that contain sensitive data and systems.
- **Threat monitoring** — Deploy advanced threat detection solutions to identify and neutralize wiper malware before it executes. Continuous monitoring of network activity can help detect anomalies that could indicate an attack in progress.
- **Threat intelligence** — Understand who may be targeting your IT environment and proactively learn about their usual tactics. Prioritize security efforts and fine-tune monitoring based on tools and techniques used by the threat actors that are more likely to attack.
- **Incident response and recovery plans** — Build comprehensive incident response and recovery strategies that help ensure quick identification, containment, and mitigation of threats. Be sure to test and exercise them regularly; a well-rehearsed plan can significantly reduce downtime during an attack.
- **Immutable backups** — Since wiper malware aims to permanently delete data, maintaining immutable backups out of reach of malware is crucial. Moreover, before unfolding the attack, threat actors can poison backups to make them unrestorable, so test your backups regularly to ensure they can actually be used for recovery.

Strengthening Cyber Resilience Against Wiper Attacks

For government organizations, cyber resilience is inseparable from mission assurance. Implementing zero-trust architecture, segmenting sensitive enclaves, enforcing least-privilege access, and maintaining immutable, routinely tested backups are no longer “best practices” but essential controls for complying with FISMA, Executive Order 14028, and CISA’s Binding Operational Directives. Combine those technical safeguards with regular tabletop exercises and an incident-response playbook that spans agencies and trusted public-private partners to transform wiper malware from an existential threat into a manageable risk. The goal is clear: ensure that even under digital bombardment, the wheels of government keep turning, and citizens never lose access to the vital services they depend on.

About the Author

Dirk Schrader is Resident CISO (EMEA) and VP of Security Research at Netwrix. A 25-year veteran in IT security with certifications as CISSP (ISC²) and CISM (ISACA), he works to advance cyber resilience as a modern approach to tackling cyber threats. As the VP of Security Research, Dirk is working on focused research for specific industries like Healthcare, Energy or Finance. As the Field CISO EMEA he 'speaks the language' of Netwrix' customers & prospects to facilitate a fit for purpose solution delivery. Dirk has published numerous articles addressing cyber risk management, IT security tactics and operations, and reported hundreds of unprotected, vulnerable critical medical devices to authorities and health providers around the globe.



Dirk can be reached online at X/Twitter [@DirkSchrader](#) , [LinkedIn](#) and at the Netwrix website: www.netwrix.com.



Cloud-Native, AI-Driven, and Always-On: The Future of Firewall Security

Next Generation Firewalls Provide Viable Security for Organizations That Take Cloud-Native Architecture Seriously

By Paul Barbosa, Global Cloud Security Business Unit Leader, Check Point Software Technologies

Firewalls were built for a different world—static networks, predictable traffic, and clear perimeters.

Today's cloud-native environments are nothing like that; applications are broken into microservices, and infrastructure spins up and disappears in minutes. Sensitive data moves between systems at machine speed. And yet, too many organizations rely on security tools that can't keep up with how modern systems actually work.

The firewall has finally evolved, but only out of necessity. The newest generation isn't an appliance or virtual machine; it's cloud-native, AI-driven, and always-on. It doesn't guard a border; it lives where your workloads live. And if it's not doing that, it's irrelevant.

But an evolved firewall by itself isn't enough, and you can't secure what you can't see—that's where most organizations are still exposed.

Cloud-Native Broke the Perimeter. You're Still Relying on It

Cloud-native architecture dismantled the perimeter model. Applications are no longer isolated behind digital moats. Workloads run in containers, scale dynamically, and often communicate in mesh networks you didn't hand-configure.

Yet security strategies still default to perimeter defense. Teams drop traditional firewalls at the edge of their VPC and call it a day, and the problem is that threats don't need to come in from outside anymore—for 83% of organizations, they're already inside. Misconfigured services, excessive permissions, and exposed secrets give attackers all they need. These vulnerabilities often live in the code and configurations your developers ship every day.

A modern firewall must be able to monitor east-west traffic, enforce policies at the workload level, and adapt to the highly dynamic nature of cloud-native environments. That requires two things: deep visibility into what's running and how it behaves and intelligence to act on that data in real time.

AI in Firewalls Isn't a Feature; It's the Only Way This Works.

There's a reason AI is everywhere in security marketing—but in this case, it's justified. AI isn't just automating rule writing. It's the only realistic way to keep pace with the scale and speed of modern workloads.

An AI-driven firewall in a cloud-native world needs to do three things well:

1. **Baseline behavior dynamically:** It must understand how your workloads behave under normal conditions, not based on static signatures but by learning patterns over time. This feature is crucial when microservices scale horizontally, and traffic patterns shift constantly.
2. **Enforce policies autonomously:** Your infrastructure is elastic, and security policies must be, too. AI-driven systems can adjust rules in real time, applying least-privilege principles to traffic flows without waiting for human input.
3. **Detect anomalies fast:** AI enables real-time threat detection by analyzing huge volumes of telemetry data, spotting outliers, and taking action before an attacker can move laterally.

This point is where AI stops being hype and starts being practical. In a cloud-native environment, static rules and manual policy updates are slow and active liabilities. Just as AI enhances core firewall functionality, AI-driven web application firewalls (WAFs) can learn application-specific traffic patterns and detect anomalies that indicate complex web attacks, such as zero-day exploits or API abuse.

Visibility: The Problem You Didn't Know You Had (But Attackers Do)

And if you can't see these exposures, neither can your firewall, where most organizations get blindsided.

They deploy a firewall thinking they've secured the infrastructure, but their actual attack surface lives in code repositories, CI/CD pipelines, and transient environments.

Modern firewalls can integrate with runtime environments, but that's not enough. You need visibility into the full software development lifecycle (SDLC), including proactive secrets detection, monitoring of build artifacts, and continuous verification.

Always-On Firewalls Demand Continuous Data—and DevOps Control

An always-on firewall isn't about passive 24/7 monitoring. It's about automated enforcement that continuously adapts to dynamic environments—no human intervention, no waiting for incident response.

But this only works if the firewall has continuous access to relevant data:

- **Workload telemetry** from ephemeral containers and microservices that spin up and down in seconds.
- **Signals from your CI/CD pipeline**, including version control, builds, and infrastructure provisioning events.
- **Identity-aware context**, tying policies to service accounts, IAM roles, and third-party integrations—not just IP addresses.
- Without this data, even an AI-driven firewall cannot effectively block unauthorized connections, enforce least-privilege access, or isolate threats. DevOps teams need to treat firewalls as part of their infrastructure—not as an external layer managed by security teams:
- **Define firewall policy as code**, version-controlled and deployed through automated pipelines.
- **Automate secrets detection** across code commits, build artifacts, and logs to prevent credential leaks before they become attack vectors.
- **Expose firewall telemetry directly to developers**, shortening feedback loops and enabling fast remediation when workloads violate policy.

Teams that embed them into their deployment workflows build more secure, scalable, resilient systems without trading off speed.

This Is the New Baseline, not a Future Vision


Cloud-native, AI-driven, always-on firewalls are already in production in forward-thinking teams. They're not a trend but the minimum viable security for organizations that take cloud-native architecture seriously.

Firewalls can't do it alone. Without visibility into your code, pipelines, and secrets, even the smartest AI can't protect you. Teams that understand this shift are already hardening their pipelines, surfacing hidden risks, and giving their security systems the data they need to work. Everyone else? Still pretending that north-south traffic is the problem.

About the Author

Paul Barbosa leads the Global Cloud Security Business Unit at Check Point, overseeing Sales, R&D, Marketing, Customer Success, and Engineering. With over 23 years in cybersecurity and communications, Paul brings deep expertise across enterprise, federal, and commercial sectors. A former U.S. Marine and security architect, he's held leadership roles at Skyhigh Security, Cisco, and Cloudflare. His leadership mantra, "Brilliant in the Basics," reflects his commitment to operational excellence and tactical precision.





Man, Machine and Malware

Roadmap of AI - Crunching Cyberattack Data to Bolstering Shields

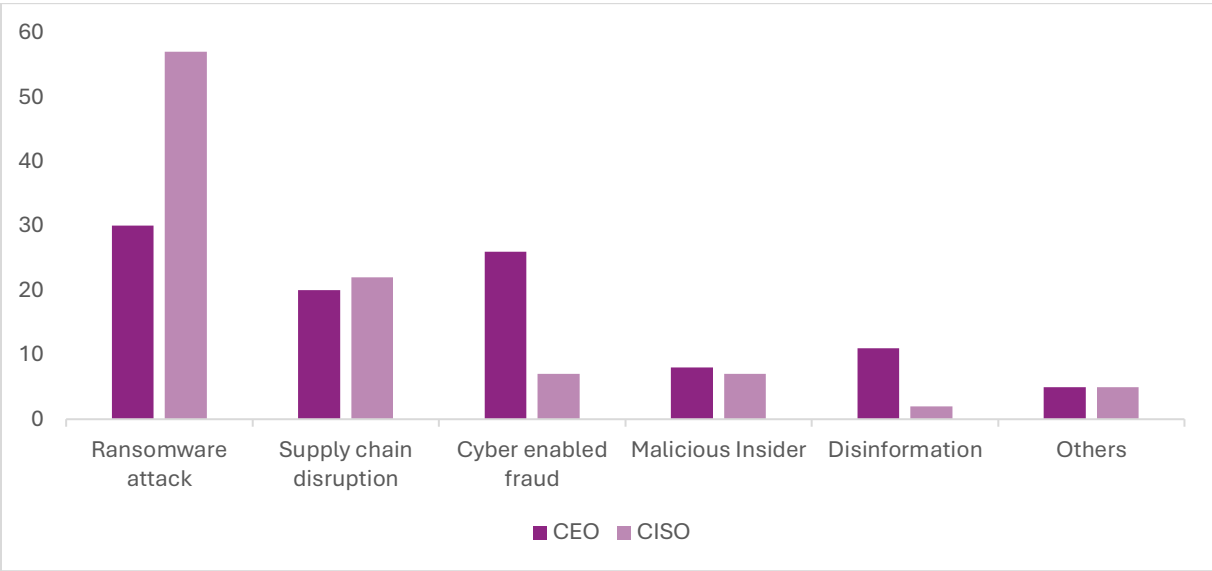
By Marlene Francis, Research Analyst, IndustryARC

With the rise of ransomware, phishing, zero-day exploits and other cyberthreats, organizations worldwide are confronting a cybersecurity crisis that traditional tools can no longer manage alone. The intersection of artificial intelligence (AI) and cybersecurity is nothing new. The cybersecurity realm has witnessed a paradigm shift since its inception in 1980s when Microsoft security team opted for rule-based systems, which gradually evolved to the current ML based adaptive security featuring baseline anomalies & behavior detection across networks. Today, innovation continues with the rise of generative AI (GenAI). This leap enables cybersecurity professionals to engage with systems through natural language and unlock deeper, more intuitive threat insights without requiring technical query syntax.

Today, AI is being integrated into Security Operations Centers (SOCs) to power behavioral analytics, identity verification, malware detection and automated response. The average large organization handles over 11,000 security alerts per day, according to IBM, which can overwhelm even experienced cybersecurity teams. The sheer volume and complexity of attacks shows the limitation of traditional monitoring. AI can sift through large data volumes to detect threats, flag anomalies and respond in real time. The World Economic Forum's 2 Global Cybersecurity Outlook 2025 shows that both CEOs and

CISOs place a high priority on AI and machine learning and underscoring their shared belief that intelligent automation is key to future-ready cyber defense.

Organizational Cyber Risk- CEO & CISO Perspective



Source: World Economic Forum

The complexity of modern cyberattacks has outpaced the capacity of traditional cyber defense

At present, cyberattacks are becoming sophisticated with [advanced persistent threats \(APTs\)](#), zero day attacks and supply chain attacks. For instance, Cisco’s 2024 Security Report states that over 3.4 billion phishing emails are sent every day. In between December 21, 2023, and July 5, 2024, Darktrace / EMAIL detected 17.8 million phishing emails across the fleet with 62% of these phishing emails successfully bypassing Domain-based Message Authentication, Reporting and Conformance (DMARC) verification checks. Google’s Threat Analysis Group (TAG) and Mandiant teams observed 97 zero-day vulnerabilities exploited in the wild in 2023, a 56% increase y-o-y. Many of these attacks are highly targeted and difficult to detect using traditional approaches. According to a 2025 IBM X-Force report, AI bots are capable of independently scanning, learning and launching attacks in real-time.

Cybersecurity experts need to sift through thousands of alerts daily. Small organizations, on an average, use between 15 and 20 cybersecurity tools, medium-sized businesses use 50 to 60 tools and large organizations use over 130 tools on average, as per figures from 2019 RSA Conference. This often leads to alert fatigue and fragmented visibility. Moreover, according to the 2024 ISC2 Cybersecurity Workforce Study, 67% of respondents stated they had faced a staffing shortage which further reinforces the fact that traditional defenses are no longer sufficient. This poses substantial consequences with respondents identifying workforce shortages as their primary challenge and anticipate that this issue will remain a major concern in the future.


AI doesn't just detect threats; it learns from them in real-time

Despite upfront investment, AI driven cybersecurity often reduces long term breach cost, labor hours and penalties. IBM's 2024 Cost of a Data Breach Report found that the average cost of data breach reached \$4.4 million in 2024 and a key finding was that use of AI and automation, on average, reduced breach costs by \$2.2 million. AI can analyze millions of datasets in real time and can detect anomalies faster. It can operate round the clock and continuously scan for threats across devices, networks and cloud services. Machine learning modules can detect subtle behavioral shifts or unseen attack patterns that traditional tools might miss. In short, AI augments cybersecurity teams by providing greater coverage, speed and context.

In a world where identity is the new perimeter, AI is the gatekeeper

Five critical areas where AI delivers the most impact are threat detection and anomaly monitoring, [Identity Access Management \(IAM\)](#), malware detection, user behavior analytics (UBA) and automated response. AI-powered systems monitor network traffic and system behavior to detect anomalies that may be an indication of malicious activity such a sudden spike in outbound traffic from a server during off hours, a user attempting unauthorized access and unusual login patterns. By incorporating ML, these systems continuously adapt to threats and anomalies to reduce false positives. Furthermore, in a world where hybrid and remote work is becoming more common, identity is the new perimeter in cybersecurity. According to IBM X-Force 2025 Threat Intelligence Index, identity-based attacks make up 30% of total intrusions. AI is reshaping IAM through adaptive authentication and biometric verification, with large language models (LLMs) capable of continuously learning user behavior and dynamically adjusting authentication and authorization protocols in real time. Traditional antivirus depends on predefined signatures to identify known threats. On the other hand, AI driven malware detection uses behavioral analysis to detect suspicious activity. As per Watchguard Threat Lab Q2 2024, 46% of malware deployed techniques that were specifically designed to evade signature-based tools. Additionally, AI enables the implementation of UBA which tracks how individuals typically behave across systems and then flags anomalies that indicates compromised accounts or insider threats. Finally, AI also powers [Security Orchestration](#), Automation and Response (SOAR) tools. By reducing mean time to detect (MTTD) and mean time to respond (MTTR), AI driven cybersecurity minimizes damage and lowers operational burden on cybersecurity experts.

GenAI Use by Cybersecurity Teams




Source: 2024 ISC2 Cybersecurity Workforce Study

Despite its advantages, AI also introduces its own set of challenges. Gen AI is a double-edged sword. As per the 2024 ISC2 Cybersecurity Workforce Study, 45% of cybersecurity teams implemented Gen AI into their teams' tools to bridge skills gaps, improve threat detection and provide vast benefits to cybersecurity. Over half have already faced data privacy and security concerns due to organizational adoption of Gen AI. As the adoption of Gen AI in cybersecurity continues to grow, cybersecurity professionals recognize the need for a formal strategy and regulations to govern its safe and responsible use.

Although 90% of respondents have some policies related to Gen AI, 65% say their organization needs to implement more regulations on the safe use of Gen AI. Another significant issue is adversarial attacks, where attackers manipulate inputs in subtle ways to mislead the system. For instance, adding seemingly benign data to a file or tweaking user behavior can confuse AI into classifying threats as safe. In fact, according to the MIT AI Risk Registry, there are potentially over 700 risks that need to be mitigated with the use of Ge AI. NIST has worked in collaboration with the private and public sectors to develop a framework to manage risks to individuals, organizations and society associated with AI and is recommending a formal AI Risk Management Framework as a standard.

Policy landscape for GenAI



Source: 2024 ISC2 Cybersecurity Workforce Study

GenAI is a cybersecurity's double-edged sword- promising defense yet empowering offence

Leading cybersecurity vendors such as Darktrace, IBM and CrowdStrike illustrate how AI is being used effectively in cybersecurity. Darktrace uses self-learning AI that is designed to autonomously adapt and detect emerging threats that traditional tools often overlook. This system is further enhanced by a multi-modal AI architecture that incorporates large language models (LLMs), GenAI techniques and supervised learning algorithms to reinforce its ability to deliver real-time, context-aware cyber threat identification. IBM Watson applies NLP (natural language processing) and machine learning to analyze massive volumes of data and provide threat context to SOC analysts. CrowdStrike Falcon uses AI to detect zero-day exploits. CrowdStrike claims its AI driven threat intelligence reduces dwell time to under 10 minutes. These three are just a few examples of several AI tools that have enabled a shift from reactive to proactive security.

Even attackers are using AI, the battle is algorithm vs algorithm

Attackers are also becoming more sophisticated by using AI tools to create deepfake phishing scams, autogenerate malware variants and mimic legitimate user behavior. As the attack surface and threats become more sophisticated, defenders are also advancing their use of AI which makes the future of cybersecurity dynamic, predictive and autonomous. The State of AI Cybersecurity Report 2025 by Darktrace found that 95% agree that AI-powered cybersecurity solutions significantly improve the speed and efficiency of prevention, detection, response and recovery. Rather than just responding to incidents, AI is evolving to predict them. By analyzing global threat patterns, vulnerability disclosures, user behaviour and patterns and geopolitical threats, predictive AI can forecast which regions are likely to be

a target, identify vulnerabilities before they are exploited and recommend pre-emptive policy changes. This allows organizations to take a proactive stance towards cybersecurity.

Countries like the United States and Singapore are actively integrating AI into their cybersecurity frameworks to bolster national resilience against evolving cyber threats. In October 2024, the White House released its National Security Memorandum on Artificial Intelligence, establishing a strategic framework for leveraging AI to advance U.S. national security objectives. To capitalize on AI's potential in strengthening both defensive and offensive cyber operations, the memorandum directs the Department of Energy (DOE) to initiate a pilot program. This initiative seeks to advance AI-driven threat detection and operational response capabilities against adversarial actors. Singapore's approach is anchored in its National AI Strategy 2.0 which includes over S\$1 billion in AI investments and frameworks like the Model AI Governance Framework for Generative AI.


The future will also see collaborative AI as organizations increasingly participate in shared [threat intelligence](#). In addition, the increase in cyber espionage particularly targeting critical infrastructure and financial institutions underscores the demand for robust intelligence frameworks. As per a white paper from Symantec, ransomware operators claimed responsibility for 1,312 incidents in Q4 2024, marking an increase from 1,255 in Q3 2024 and up from 1,179 during Q4 2023. Legacy security tools, are increasingly inadequate against modern threat vectors. In response, organizations are prioritizing threat intelligence as a strategic defense layer to enable proactive detection, contextual analysis and mitigation of cyber risks before they materialize. According to research by Cyware, 92% of cybersecurity professionals emphasized that collaborative threat intelligence sharing is highly valuable, which reinforces its role in today's dynamic threat landscape.

The future of cybersecurity is not man or machine. It's man with machine

Artificial Intelligence is already embedded into cybersecurity and will continue to play a crucial role in defense tactics. The 2024 ISC2 Cybersecurity Workforce Study reports that AI/ML has entered the top 5 in-demand skills for cybersecurity professionals. Cybersecurity professionals ranked automation (42%) and AI (36%) as the two technologies that will have the greatest impact on their ability to secure their organization. With threat actors themselves leveraging AI to develop more sophisticated attacks, the integration of AI in cybersecurity offers organizations the ability to scale their cybersecurity efforts while reducing response times and human workload. In conclusion, the most effective cyber strategies will combine the speed and precision of AI with the contextual understanding and ethical judgements of human experts. AI shouldn't be viewed as a replacement for the human cybersecurity workforce, but rather a force multiplier for cybersecurity teams. Let AI be a powerful ally that becomes the cornerstone of resilient and intelligent cyber defense systems.

About the Author

Marlene Francis is a Research Analyst (Level-2) at IndustryARC, where she leads content strategy and development across key technology verticals. With over four years of experience, she has contributed to market research and thought leadership. At IndustryARC, Marlene plays a pivotal role in translating complex insights into accessible, high-impact narratives for business leaders and industry stakeholders. She is passionate about bridging the gap between technical research and real-world decision-making. Marlene can be reached at marlene.francis@industryarc.com and at the company website: <https://www.industryarc.com/>





Threat Intelligence in the Heat of Cyber Warfare

By Andrew Martin, CEO at DynaRisk

When conflict erupts on the world stage, our attention is drawn to the more tangible threats, be it missiles in the sky or troops on the ground. Yet behind these visible events, bubbling under the surface lies another pressing danger of war, often fought in silence and in the shadows. On this particular battlefield, weapons are lines of code, the terrain is digital infrastructure, and the damage, though less immediately visible, can still be devastating.

Threat actors aligned with nation states are increasingly targeting not just military systems, but the civilian networks that keep society running. Hospitals, energy providers, transportation hubs and communication systems are all potential victims. Disabling these assets can slow military responses, hinder emergency services and disrupt public order without the need to fire a single bullet.

The evolution of cyber tactics

Modern cyber warfare is no longer confined to stealing intelligence or quietly monitoring adversaries. Today's attackers seek to exploit software flaws, manipulate supply chains, and deploy highly targeted social engineering campaigns. The interconnected nature of critical infrastructure means that a breach in one area can quickly ripple through others, debilitating crucial systems at the click of a button. Just

consider for a moment the impact that shutting down an entire food supply chain or logistics system could have if it prevented military resupplies during a time of war.

The fact is a single phishing email has the potential to open the door to a national crisis. The threat of cyber warfare is real.

Global fallout beyond the battlefield

The most concerning feature of cyber warfare is just how quickly malware, ransomware and other threat types can spread across borders. The technical nature of such attacks means that they're not limited by the laws of time and speed that traditional warfare methods are and can therefore impact new locations and whole nations almost immediately. A pressing example of this was the NotPetya incident of 2017, intended to disrupt systems in Ukraine, quickly spread globally, paralysing shipping companies, manufacturers and retailers in dozens of countries.

This reality means no organisation can consider itself outside the threat radius when the conflict is being fought on the cyber battlefield. Whilst not the primary targets, businesses large and small may find themselves ensnared in a large-scale incident simply by being connected to global supply chains.

The consumer as a target

Cyber warfare also reaches into the daily lives of ordinary citizens. Disinformation campaigns are often used to manipulate perceptions, polarise communities and undermine trust in public institutions. As part of these tactics, phishing scams can masquerade as urgent messages from government agencies, ultimately luring individuals into revealing personal information or installing malicious software.

Even stolen personal data has strategic value. When aggregated, it can help adversaries identify and manipulate individuals who hold influence over security, politics or public opinion. This makes consumers an integral part of the threat landscape, whether they realise it or not and acts as a foot in the door for these malicious threat actors.

The insurance challenge

All these challenges have considerable knock-on effects for cyber insurance. While traditional policies are often designed around isolated, predictable risks, cyber warfare introduces new systemic risks in the form of attacks that cause simultaneous, widespread damage across multiple industries and geographies.

This raises complex questions about coverage. Many policies exclude acts of war, leading to disputes about whether a given incident falls under that definition. Insurers may struggle to meet claims in the face of a mass-scale cyber event, leaving businesses unexpectedly exposed at the moment they most need financial support.

Intelligence as a strategic imperative

In this environment, threat intelligence becomes more than a technical exercise and becomes a strategic necessity. The most resilient organisations will be those that can detect emerging threats early, interpret them in the context of the wider geopolitical landscape, and take pre-emptive action.

This requires more than monitoring internal systems for signs of compromise. It means tracking the activity of threat actors, understanding their motivations, and recognising when geopolitical shifts increase the likelihood of attack. It's critical that intelligence is translated into decisive action, be it patching vulnerabilities, strengthening authentication, or rehearsing incident response.


The landscape on which Geopolitical conflict is fought is changing rapidly. No longer fought solely on land, sea and air, it has unfolded into data centres, on social media platforms, and across the networks that underpin our economies. The question is not whether cyber warfare will touch your organisation, but more likely when. What can be controlled is how ready you will be when it does.

Organisations that see threat intelligence as a core leadership priority, rather than a background IT function, will be best placed to weather the storms of this new era where the stakes are no longer limited to shareholder value or brand reputation. In a world where critical infrastructure, economic stability and public trust can be undermined with a malicious click, the decisions leaders make today will determine their resilience in the battles to come.

About the Author

Andrew Martin is the founder and CEO of DynaRisk, a London-based cyber risk management firm he launched in 2016 after recognizing that individuals, families, and SMEs lacked simplified cybersecurity tools powered by strong threat intelligence. Andrew can be reached at our company website www.dynarisk.com





Deepfakes, Synthetic Media, and Digital Trust: The Cybersecurity Implications of Deepfake Technology and Methods for Detection and Mitigation

By Joe Guerra, M.S., Software Engineering, CASP+, CCSP, FedITC,LLC

Introduction

The proliferation of deepfake technology, synthetic media generated using advanced artificial intelligence techniques, has emerged as a significant challenge to digital trust in the modern era. Deepfakes, a term derived from “deep learning” and “fake,” involve the manipulation or creation of video, audio, or images to convincingly mimic real individuals and events. Fueled by rapid advancements in generative AI, especially Generative Adversarial Networks (GANs), deepfakes have progressed from a curiosity in online forums to a tool capable of undermining personal, organizational, and societal security. This article explores how deepfake technology operates, the implications of its misuse in cybersecurity, and ongoing advancements in detection and mitigation.

The Technology Behind Deepfakes

Deepfakes are primarily created using machine learning algorithms, especially GANs. A GAN consists of two neural networks, the generator and the discriminator, that work together to create data indistinguishable from real-world examples. The generator crafts synthetic media, while the discriminator evaluates their authenticity, prompting iterative improvements in realism. As a result, GANs can produce highly convincing images, videos, and voices.

The accessibility of deepfake creation tools has expanded dramatically. Open-source libraries and commercial platforms now allow individuals with modest technical skills to produce deepfakes. Video and audio impersonation can be achieved with relatively small datasets, sometimes just a few minutes of video or audio of the target. The rise of deepfake-as-a-service platforms means attackers can now outsource the creation of synthetic content, lowering barriers for cybercriminals and increasing the proliferation of malicious synthetic media.

Real-world Threats and High-profile Incidents

Deepfakes have already played a role in several documented cyberattacks. In 2019, a deepfake audio attack successfully impersonated a CEO's voice to trick a subordinate into transferring €220,000 to a fraudulent account. Similar incidents have targeted companies, government agencies, and individuals, with attackers using synthetic voices or faces to bypass security controls or manipulate victims.

Cybercriminals are increasingly leveraging deepfakes for spear-phishing campaigns and business email compromise (BEC) scams. Adversaries can use AI-generated audio or video to impersonate executives or trusted officials, lending authenticity to fraudulent requests. Political deepfakes have been deployed to spread misinformation or undermine public trust during elections, fueling disinformation campaigns that can influence opinions and destabilize societies.

Moreover, deepfakes are used for extortion, harassment, and blackmail. Malicious parties can fabricate compromising images or videos, threatening personal reputations and causing emotional and psychological harm. As the technology advances, it becomes more challenging for victims, and even experts, to distinguish real from fake, escalating the overall risk landscape.

Impact on Digital Trust and Society

At an individual level, deepfakes pose severe risks to privacy, reputation, and financial security. Synthetic media can be used for identity theft, impersonation, and personal attacks that leave lasting damage on victims' lives. For organizations, deepfakes open the door to social engineering attacks, corporate espionage, fraud, and significant brand damage. A convincing deepfake can trick employees, investors, or customers, resulting in both direct financial loss and long-term erosion of trust.

Democracies face even broader implications. Deepfakes can spread false information on social media, fueling disinformation campaigns that erode trust in public institutions and the media. Fabricated videos

and audio used in election cycles can tarnish candidates, misinform voters, and undermine the legitimacy of democratic processes.

Perhaps most fundamentally, the rise of synthetic media challenges the principle of “seeing is believing.” As the line between reality and fabrication blurs, the authenticity of digital communications, online evidence, and documentation becomes questionable, threatening the fabric of digital trust society relies upon.

Detection Techniques and Their Limitations

AI/ML-based Detectors

Detection of deepfakes typically relies on artificial intelligence and machine learning models trained to identify anomalies characteristic of synthetic content. These detectors analyze artifacts such as inconsistencies in facial movements, blinking rates, light reflections, or audio-visual mismatches. Some algorithms can detect subtle unnatural features left behind by GANs or identify statistical differences between real and fake samples.

Biometric and Watermarking Solutions

Biometric analysis focuses on physiological traits, like heart rate inferred from facial video or micro-expressions difficult for GANs to reproduce. Meanwhile, digital watermarking involves embedding invisible markers or digital signatures in authentic media, which can later be checked to verify integrity.

Crowd-sourced and Manual Verification

Manual verification by trained professionals, journalists, or fact-checkers remains a valuable approach, especially when automated tools produce inconclusive results. Crowd-sourcing platforms and community-driven services can rapidly vet viral content, although such efforts are labor-intensive and not always timely.

Blockchain and Authenticity Tags

Blockchain-based systems and tamper-evident metadata can offer methods for recording and tracing the provenance of digital media. By establishing audit trails and issuing authenticity tags at the point of creation, these systems help verify that media has not been altered.

Limitations

No detection method is foolproof. As deepfake generation methods improve, so do their abilities to evade AI-driven detectors. The ongoing arms race between creators and defenders of synthetic media means detection algorithms require continuous updating and retraining as new attack methods emerge. Moreover, detection is further complicated when deepfakes are of low resolution, fleeting in nature (e.g., live streams), or specifically tailored to defeat known detection models.

Mitigation Strategies and Best Practices

Technical Safeguards

Organizations are increasingly deploying deepfake detection APIs and integrating provenance tools within security infrastructures. Real-time content monitoring, digital watermarking, and media integrity verification solutions can be incorporated into platforms where media is uploaded or shared.

Authentication protocols should be enhanced: for sensitive transactions or communications, relying solely on voice or video authentication is becoming inadequate. Implementing multifactor authentication, such as biometrics, tokens, or behavioral analytics, helps mitigate risks posed by deepfake-enabled impersonation.

Policy and Organizational Responses

Employee training and awareness are essential. Security teams should educate staff on the existence and detection of deepfakes, emphasizing skepticism toward urgent or unusual requests delivered through audio, video, or messaging apps. Incident response playbooks must now include protocols for investigating and responding to deepfake-enabled threats.

Collaboration within industries and across public and private sectors is vital in developing threat intelligence, sharing best practices, and standardizing verification methods. Legal mandates requiring clear labeling of synthetic media, accountability for creators, and timely reporting of deepfake attacks are emerging as critical tools in the policy arsenal.

Legal and Regulatory Developments

Legislation addressing deepfakes is evolving. Some jurisdictions impose criminal penalties for malicious deepfake use, mandate disclosures for synthetic content, or empower regulators to oversee social media compliance. However, legal frameworks must balance protection from harm with rights to free expression and innovation.

Future Outlook

The realism and ease of access to deepfake technology are expected to increase, expanding the threat landscape and complicating detection efforts. Upcoming advances in generative AI will allow ever more convincing impersonations, making non-technical users potential creators and victims alike.

Addressing these evolving threats requires a multidisciplinary approach: technical innovation in detection tools, legal reforms, widespread awareness, and a rethinking of how digital trust is established and verified in an age of synthetic media. The cybersecurity community must collaborate with technologists, policymakers, and society at large to keep pace with new risks.

Conclusion

Deepfakes and synthetic media challenge fundamental tenets of digital trust, exposing individuals, organizations, and democratic institutions to unprecedented risks. As the technology progresses, so too must the strategies for detection, mitigation, and public awareness. Balancing innovation with resilience requires proactive investment, multidisciplinary research, and collaboration across sectors. By remaining vigilant, adopting robust safeguards, and fostering digital literacy, society can confront the deepfake threat and preserve the foundations of trust in our increasingly synthetic digital world.

References

1. Westerlund, M. (2019). The emergence of deepfake technology: A review. *Technology Innovation Management Review*, 9(11), 39-52.
2. Mirsky, Y., & Lee, W. (2021). The creation and detection of deepfakes: A survey. *ACM Computing Surveys*, 54(1), 1-41.
3. Tidy, J. (2019, August 30). Energy boss scammed out of millions in AI voice fraud. *BBC News*.
4. Chesney, R., & Citron, D. K. (2019). Deepfakes and the new disinformation war. *Foreign Affairs*, 98(1), 147-155.
5. Paris, B., & Donovan, J. (2019). Deepfakes and cheap fakes: The manipulation of audio and visual evidence. *Data & Society*.
6. Vaccari, C., & Chadwick, A. (2020). Deepfakes and disinformation: Exploring the impact of synthetic political video on deception, uncertainty, and trust in news. *Social Media + Society*, 6(1), 1-13.
7. Verdoliva, L. (2020). Media forensics and deepfakes: An overview. *IEEE Journal of Selected Topics in Signal Processing*, 14(5), 910-932.
8. Kietzmann, J., Lee, L. W., McCarthy, I. P., & Kietzmann, T. C. (2020). Deepfakes: Trick or treat? *Business Horizons*, 63(2), 135-146.
9. Nguyen, T. T., Nguyen, C. M., Nguyen, D. T., Nguyen, D. T., & Nahavandi, S. (2022). Deep learning for deepfakes creation and detection: A survey. *Computer Vision and Image Understanding*, 214, 103329.

About the Author

Joe Guerra, M.S., Software Engineering, CASP+, CCSP, RMF ISSO/ISSM Instructor, FedITC, LLC. San Antonio, Texas (Lackland AFB)

He is an experienced computer science and cybersecurity educator with over 20 years of expertise. He spent 12 years teaching science, Information Technology, and Computer Science at the high school level, shaping young minds and inspiring the next generation of technology professionals. His deep knowledge and passion for the field paved the way to higher education. Joe holds three Masters: Master's degree in Information Systems Security, Software Engineering, and Instructional Technology, and is certified in CompTIA Network+, Security+, CySA+, and CASP+, as well as CCSP by ISC2.




For the past 10 years, Joe has been an esteemed adjunct instructor at ECPI University, the University of the Incarnate Word, and Hallmark University. He has taught a wide range of courses, including Security Assessment and Testing, Identity and Access Management, Linux operating systems, and programming languages such as Java, C, Python, C#, and PowerShell. His diverse skills also encompass networking, cybersecurity, Cisco systems, Hacking and Countermeasures, and Secure Software Design.

A highlight of Jose's career was his 2019–2023 role teaching Air Force cyber capability developers, where he focused on developing offensive and defensive software tools, making significant contributions to cybersecurity warfare and national defense.

In addition to his technical teaching, Joe specializes in training cyber leadership personnel, including Information System Security Officers (ISSOs) and Information System Security Managers (ISSMs), in the Risk Management Framework (RMF) process. He equips these cyber professionals with the knowledge and practical skills required to navigate complex regulatory environments, ensure compliance with federal standards, and implement robust security controls. Joe's instruction emphasizes real-world application of RMF, fostering a deep understanding of risk assessment, security authorization, and continuous monitoring. His approach prepares ISSOs and ISSMs to lead cybersecurity initiatives, manage enterprise risk, and uphold the highest standards of information assurance within their organizations.

Joe's dedication to education, hands-on expertise, and leadership in both technical and managerial aspects of cybersecurity make him a trusted mentor and resource for students and professionals alike. Joe can be reached online at jguerra@Feditc.com, joe.guerra@afcsu.org, mrjoeguerra@outlook.com)and at our company website <https://feditc.com/>



How Red Teams are Reinventing Cybersecurity for the Age of AI

As generative AI systems permeate critical infrastructure, red teams must adapt their tactics to confront unpredictable behaviors, data leakage, and multi-modal vulnerabilities.

Red teaming has long served as a cornerstone of cybersecurity, probing networks and platforms for flaws before attackers can exploit them. Now, these cyber professionals are turning those same tactics to artificial intelligence systems to help ensure trust and security.

However, testing AI requires different evaluation tactics. Conventional IT systems behave predictably. Not generative AI. Its non-deterministic nature can produce different results from the same prompt. It also lacks fixed rules, making it harder for red teams to test and secure consistently.

Ironically, the very unpredictability that complicates red teaming also makes it indispensable, especially as AI plays a growing role in essential societal functions and public safety. Let's take a closer look at the main risks tied to GenAI and the emerging strategies red teams can use to strengthen its security and reliability.

Why Red-Teaming is Mission-Critical for AI

The primary reason for red team exercises is a lack of trust in a system's security and accuracy. Organizations with questions about their AI unleash red teams to "break" the system.

Typically, the teams look for two things. The first is data leakage, where sensitive data is input into GenAI systems, then extracted through adversarial techniques. The second is erroneous outputs, in which Large Language Models (LLMs) produce inaccurate or misleading information.

Most GenAI tools have guardrails in place to prevent users from getting hold of certain information, such as proprietary, dangerous, or private content like medical records or blueprints of critical infrastructure. They limit the types of questions and prompts the systems allow. However, attackers employ various tactics to manipulate systems, including using clever wording to trick LLMs into ignoring their presets, injecting malicious instructions through lengthy or complex prompts, and more.

Successful attempts at AI manipulation can compromise emergency response effectiveness. Think of a situation in which emergency medical technicians or law enforcement officers are using AI to determine the origin point of a 911 call. Attackers could inject erroneous information into an LLM that could significantly delay response times.

Three Innovative Red Teaming Strategies for AI

There are several innovative strategies red teams can adopt to match attackers' increasingly sophisticated methods.

Prompt attack red teaming

Red teams can create prompts designed to bypass the AI's guardrails and rephrase them multiple times using different words, phrases, idioms, and typos. They can also hide malicious commands within long strings of text to test whether the AI accepts them, or ask the AI to provide private data or deliver inaccurate results.

Red teams can also train internal AI systems to recognize patterns typical of prompt attacks and alert security teams when suspicious activity is detected. This type of counter-processing acts as a safety net in case an adversary bypasses the red team's defenses.

Multi-modality red teaming

Many organizations deploy AI for object detection, facial recognition, vehicle identification, and voice commands, or to determine the meaning behind a scent or sound. Adversaries with the right skills and technology can easily exploit these multi-modal channels and take advantage of the expanded attack surface. For example, a hacker could subtly adjust the visual patterns of street signs, causing them to disappear, making it impossible for autonomous vehicles to navigate safely.

Organizations should conduct adversarial red teaming to mitigate these challenges. Adversarial red teaming involves deliberately trying to undermine multi-modal systems by feeding them deceptive inputs.

For instance, red teams could make imperceptible changes to imagery, input a combination of audio, visual, and text prompts, or intentionally poison model data with inaccurate information about objects.

Digital twin red teaming

AI is constantly interacting with and learning from the world at large; its interface is wide open. Therefore, red teams must be able to replicate real-world scenarios that could take place if their AI is compromised.

Digital twins are virtual representations of real-world systems, processes, or environments. Red teams can utilize digital twins to simulate "what if" scenarios, depicting the impact of a malfunctioning AI on an actual process—for example, what would happen if an attacker were able to manipulate a city's emergency response system.


Red teams can create digital twin versions of their AI systems—realistic replicas that mirror the model, logic, and how the system interacts with its environment. These twins let teams safely experiment with attacks like prompt hacking or tampering to see how the real system might react, without putting anything critical at risk.

Blending Traditional and Modern Testing for Maximum Security and Reliability


While red teaming must evolve to meet the fluid nature of AI, traditional testing methods, such as penetration testing and monitoring attack surface analytics, remain essential. They, along with the techniques outlined here, create a powerful, layered approach to ensuring secure and reliable AI, which will help build public trust in the technology.

About the Author

Darren Pulsipher is the chief solutions architect for the public sector at Intel. He works directly with governments (federal, state, and local) and enterprise organizations such as IBM, GE, and Toyota to help them modernize their IT organizations. Through several executive and management positions (CIO, director of engineering), Darren has developed a unique ability to bring technology, people, and processes together to provide real transformational change to organizations. He focuses on data transformation through data architecture, workload migration, cloud-native application development, service orchestration, and multi-hybrid cloud data center architectures. His research has resulted in eight patents in cloud and grid computing architectures, helping companies decrease product development lifecycle time through build, test, and deployment optimization, virtualization, and containerization. Darren shares his passion for digital transformation on his weekly podcast "Embracing Digital Transformation." He is a published author with three books on technology and technology management and over 100 articles published in various industry trade publications.



More information on his work can be found at [Dr. Darren Speaks](#)



Inside the Mind of a Threat Actor: What CISOs Must Learn Before the Next Breach

Using Attacker Psychology to Strengthen Enterprise Defenses

By Ahmed Awad (aka nullc0d3), Senior Cyber Threat Intelligence Analyst, Author & Educator

Cybersecurity isn't a game of defense—it's a game of anticipation. Yet too many CISOs and security leaders still think in terms of controls, compliance, and detection thresholds. Meanwhile, the adversaries think like hunters. They exploit mindset gaps as much as technical ones. To close the breach gap, CISOs must begin thinking like attackers.

Over my 20+ years as a Cyber Threat Intelligence Analyst and Red Team strategist, I've learned one truth that governs this field: underestimating attacker psychology is the biggest blind spot in most enterprise security programs.

Threat Actors Aren't Just Technicians—They're Strategists

Whether you're facing an APT backed by a nation-state or a skilled ransomware affiliate group, their success rarely hinges on zero-days. It hinges on knowing how defenders think—and then staying one step ahead.

Attackers rely on reconnaissance, deception, and lateral movement because they understand the human patterns within enterprise security:

- SOC fatigue = more alert noise = easier evasion
- Over-reliance on signatures = blind spots to behavioral anomalies
- Rigid playbooks = predictable response windows

Their mindset is: "How can I live inside the environment without triggering alarm bells?" That's not hacking—it's infiltration psychology.

Case Study: Operation Quiet Wolf

In one red team engagement I led, we simulated a persistent adversary against a financial firm. Rather than launching a brute-force phishing campaign, we weaponized patience. We spent two weeks profiling helpdesk behavior, building spoofed identities that mimicked internal contractors. The initial access came from a Slack impersonation, not a malicious payload.

We bypassed EDR not by disabling it, but by using signed binaries and trusted paths. We didn't trip the alarms—because we thought like the blue team and danced around its visibility.

Lesson: Tools evolve, but attacker psychology—the hunger to blend in—remains consistent. And most defenders aren't trained to anticipate that level of discipline.

What CISOs Must Adopt from Threat Actors

1. Asymmetric Thinking: Attackers look for what's not expected. CISOs must challenge their teams to threat-model their own systems as if they were the adversary.
2. Deception as Defense: Honeypots, fake credentials, and traps aren't optional anymore—they're necessary to increase attacker cost.
3. Live-Fire Testing: Annual pen-tests are outdated. Simulate persistent threats using red/purple teaming with adversary emulation frameworks like MITRE ATT&CK.
4. Emotional Intelligence: Understand that threat actors often play on psychology—urgency, trust, and routine. Defensive awareness training must address human behavior, not just phishing.

The Red Team Mindset: CISO Readiness Checklist

- [] Does our SOC recognize low-and-slow TTPs?
- [] Do we monitor for identity-based anomalies (not just malware)?
- [] Can we detect lateral movement without relying solely on EDR?
- [] Are we logging identity provider events?
- [] Do we run regular threat emulation scenarios?
- [] Is threat intel actionable, not just reactive?

Final Thoughts

Defenders don't need to become attackers. But they must understand their psychology. The battlefield has shifted: it's no longer about building higher walls, but understanding who's trying to climb them—and how.

CISOs who embrace this mindset won't just be harder to breach. They'll be impossible to predict.

About the Author

Ahmed Awad, known online as nullc0d3, is a Senior Cyber Threat Intelligence Analyst with over 20 years of hands-on experience in offensive and defensive cybersecurity. He's the author of *Inside the Hacker Hunter's Mind* and *Inside the Hacker Hunter's Toolkit*. Ahmed has trained red and blue teams globally, and specializes in adversary emulation, malware analysis, and cyber warfare strategy. He can be reached on LinkedIn, Twitter (@NullC0d3r), or at <https://ahmedawadnullc0d3.pro>





AI Adoption in Application Security

Application Security leaders' strategic guide on AI adoption for AppSec teams

By Kaushik Majumder, Director, Sony India Software Centre

Artificial intelligence (AI) is the latest disruptor in technology space which is making huge impact in all industry sectors & business functions. Cybersecurity leaders have also been working on adopting AI in day-to-day operations, most popular being cyber defense, but there is potential to adopt cyber security in various other domains like Third party risk management, application security, security audits etc.

This article focusses on how to leverage AI in application security lifecycle. Broadly, below are the phases of application security

- Secure code development
- Static & dynamic security testing
- Application penetration testing
- Vulnerability remediation & reporting

Most organizations will require large teams with critical skills to ensure that security is maintained at various stages of the application security lifecycle.

The challenge then becomes to demonstrate return on investment (ROI) & the constant pressure to keep these applications & platforms safe from relentless outside attacks

Leveraging AI can argue the human capability & support in enhancing operational efficiency

Below are sample use-case types which can be considered for AI integration in application security

- Summarizing guidelines, best practices, specific instructions & reference secure codes based on secure coding policies of the organization for different programming languages
- False positive removals from static and dynamic security testing tools & reports
- Integration of AI in manual Pen testing
- Leverage AI to prioritize vulnerability remediation
- Leveraging AI in test/assessment report creation

The above AI use-cases are at different level of maturity within the industry today, while there are many tools which demonstrate good success in vulnerability remediation prioritization, removal of false positives & providing secure coding references, some of the use-cases will still need more maturity like AI integration in manual pen testing

Leveraging AI in manual pen testing is an interesting topic, one that probably needs more research and maturity, however the question that everybody is asking, whether manual pen test (humans) will be replaced by AI enabled pen tests?

In penetration testing which requires critical thinking, logical co-relation & applying contextual knowledge during the testing itself, there is less probability that AI enabled pen testing will eventually replace humans. But AI will augment the capabilities of humans & support in scaling up the operations, hence it is important for pen testing teams to develop AI functional understanding of the models enabling AI pen testing & capability to operate and manage AI enabled pen testing tools/apps built on top of those models

Augmentation of the human

Pre-testing:

Thorough understanding of the application

- The most important part of conducting a successful pen test is to first understand the application, its business requirement/logic, environment to gather as much as contextual & technical/security information as possible
- AI chatbots can be used to create a structured document based on the transcript of the meeting, so that AI can quickly provide the information in order of priority in a structured and well-defined manner
- This will reduce the time of the tester to understand the application & also reduce overall response time

Development of Test cases

- Pen testers develop test cases based on the testing target, which is mostly developed by referring to various international standards like Penetration Testing Execution Standard (PTES), pen Web Application Security Project (OWASP), Open Source Security Testing Methodology Manual (OSSTMM) etc.
- This can be a cumbersome effort as multiple sources need to be referred to & test cases need to be created
- AI can be used here to develop the preliminary test cases by referring to the most critical/acceptable industry frameworks (> 2 or 3, which may be humanely difficult)
- These test cases can then be finalized for each target by a human by applying the contextual application knowledge (Business knowledge, architecture, criticality etc.)

During Testing:

Information gathering phase/reconnaissance

- Before an actual pen test is conducted, different types of open source & proprietary tools are used to automatically detect and identify vulnerabilities on the target, but the reports generated have a lot of data & noise in it
- Tools which have AI integrated in it may provide a more refined & intelligent report which may be possible for the tester to immediately start using & thereby reducing time & effort

Exploitation Phase

This is the most important & complex phase of the pen testing activity. It may require a lot of maturity of the processes & skills within the function & pen testers to start leveraging AI in this phase. Leveraging AI in this phase requires a functional understanding of LLMs or Reinforcement based learning models (Most popular models which are the base on which AI pen test tools are created), along with hands-on experience with the tool itself, so how do we achieve this?

- Develop a training plan which includes development of functional knowledge on LLMs, Reinforcement based learning models
- Develop practical use-cases & hands-on experience in working with AI enabled pen testing tools like PentestGPT, Deep exploit, etc.
- Based on the experience gathered, develop standard operating processes on how human pen testers can leverage AI in exploitation phase
- Start with simple use-case first

The benefits of AI enabled pen test vs traditional manual pen test was demonstrated in research published in "International Journal of Scientific Research in Computer Science, Engineering and Information Technology" on 12th Dec'2024

Performance Metric	Traditional Testing	AI-Powered Testing	Improvement
Testing Time	Baseline	76% Reduction	+76%
Vulnerability Detection Rate	Baseline	2.5x Increase	+150%
False Positive Rate	Baseline	62% Reduction	+62%
Resource Utilization	High	Optimized	~40% Reduction
Coverage Depth	Limited	Comprehensive	~85% Increase
Real-time Monitoring	Manual	Continuous	24/7 Coverage

Table 1: Comparison of Traditional vs. AI-Powered Penetration Testing Performance Metrics [7, 8]

Report writing:

Report writing takes lot of time for pen testers, AI can be leveraged to enable the pen testers, where it can write the basic report based on inputs shared by the pen testers, which can then be customized this will reduce the turnaround time for pen testing

In conclusion, AI implementation needs a long-term approach, & needs to be customized based on specific situation & maturity of the organization, in context of pen testing, whether to leverage AI by breaking down the entire process and developing specific use-cases or leveraging an AI tool for the end-to-end process can be decided by the organization

In both above approaches a combination of inhouse built AI tools & third-party tools may be required

References:

<https://www.eccouncil.org/cybersecurity-exchange/penetration-testing/ai-and-cybersecurity-in-penetration-testing/>

<https://ermprotect.com/blog/how-artificial-intelligence-will-drive-the-future-of-penetration-testing/>

<https://arxiv.org/abs/2308.00121>

https://www.researchgate.net/publication/387043979_Revolutionizing_Penetration_Testing_AI-Powered_Automation_for_Enterprise_Security

Note: This article is written for the sole purpose of sharing knowledge within Cybersecurity community & opinions expressed in this article is personal & not related to any organization or entity

About the Author

Kaushik Majumder is the Director, Information Security at Sony India Software Centre. He is an information Security Strategy & Operations Leader & has 18 years of experience in information security. He is a trusted advisor to CISOs & Business leaders. Kaushik's core area of expertise is Strategy, Risk Management, & setting up Security Center of Excellence from India.

Kaushik can be reached online at www.linkedin.com/in/kmajumder1982





Dormant Access and the Hidden Risk inside your IAM Program

By Durgaprasad Balakrishnan, Independent Cybersecurity Researcher and Director of Cybersecurity – Identity and Access Management at a Leading Global Fintech Company

Dormant access refers to any account or entitlement that keeps its privileges but shows no sign of use for an extended period. This can be a domain admin that has not logged on in 90 days, a Linux service account which was never used after it was created or a SaaS admin role that never calls the API it was created to manage. They are still authorized, valid and still sitting there unused.

Dormancy can be classified into the following buckets:

1. **Human accounts** - No interactive login, MFA push or token refresh for X days (often 60–90 days).
2. **Service and machine identities** - No process start, keytab request or secret rotation in Y days (commonly 120–180 days). This can also include the container image or VM which was decommissioned, but its service principal still lives in the vault.
3. **Entitlements and roles** - The account logs in, but the specific entitlement (admin role, S3 bucket policy, sudo rule) has not been exercised within the threshold.

The exact number of days to count for dormancy may vary by industry, but the underlying idea is the same. If an account or its privileges has not been touched in two or three full business cycles, it is dormant. At that point, it turns dangerous to overlook those accounts.

Why Dormancy Matters

Most breaches that you hear start with a phishing link or a zero-day. The one you rarely hear is the attacker who simply logs in with an account that no one has touched. Dormant accounts and entitlements (memberships, roles and keys) that sit untouched yet fully empowered, are the quietest and easiest way into any environment. They stick around through reorgs, mergers, cloud moves and even outlast multiple CISOs. If no one notices an account is idle, no one notices when it suddenly springs back to life, except the attacker who stole it.

Common Hiding Spots

Here are some examples of various places where dormant accounts and entitlements can be found:

- Active Directory: stale domain accounts without any last login timestamp or sometimes disabled but still has the delegation rules that grant Kerberos tickets.
- Unix and Linux: orphaned service IDs never aged out by PAM.
- SaaS platforms: admin accounts left behind by contractors who moved on months ago.
- IAM policy sprawl: nested roles in AWS IAM or Azure AD groups that belong to a project which was sunset.

What this really means is that dormancy is not a single problem, it is a combination of thousands of accounts, entitlements and policies that are scattered across multiple systems and platforms. Bringing all this information into a centralized platform, such as a modern IGA solution, connects these scattered insights and makes it easier to take timely action.

How to Detect Dormant Access

Finding dormant accounts doesn't have to be complicated, your existing tools and logs usually have all the necessary data. Here are some of the ways you can quickly identify unused access without extra overhead:

- Windows Accounts: Use your logging platform (like Splunk) to quickly surface accounts that have not logged in for a few months. Look for users with no activity for 90 days or more. This will be enough to clearly separate inactive accounts from active ones.
- Active Directory (AD): Run built-in AD reports or use powershell scripts to list users who have not logged in recently. This can be done using the in-built AD attributes (last login timestamp/ last logon). A typical window is about four months, accounts inactive that long are almost certainly

dormant. If active directory domain is integrated to an IGA system for provisioning and aggregation, the detection and remediation of these accounts can be continuous.

- Linux: Check the last-login logs on Unix and Linux systems. Any accounts that show “never logged in” or have very old last-login timestamps should raise immediate flags.
- Cloud Identities: In AWS, Azure or GCP, regularly review identity and access reports to find roles or users with no recorded activity. Most cloud providers offer simple reports showing “last used” timestamps which is perfect for quickly spotting dormant identities.

There is another category of dormant accounts in applications. Detecting dormancy for application accounts is a must but little tricky. Application accounts often have non-interactive usage as they may authenticate through APIs or backend services, not following the traditional login methods. Usage data may be buried in generic application logs or system events. Also, a single application account might support several backend services or scheduled jobs. Using targeted log aggregation, parsing and periodic re-validation of access with the application owners will help in this process.

Tracking Dormancy Cleanup

Cleaning dormant access isn't just good security hygiene, it is a concrete, measurable step toward regulatory compliance. Standards like PCI DSS, SOX, and NIST explicitly call out stale and inactive access as risks that must be managed:

- PCI DSS 4.0 (Req. 7.2.5) explicitly states that user accounts and their access levels must be reviewed regularly. Actively purging dormant accounts not only addresses this requirement directly but also gives you clear evidence to show auditors you are on top of your controls.
- SOX 404 targets unused privileged accounts as a prime example of material weaknesses. When you automate dormancy cleanups, auditors see fewer exceptions. The smaller the sample of problematic accounts, the smoother and faster is your audit.
- NIST Cybersecurity Framework v2.0 (PR.AC-1) clearly instructs organizations to manage identities by disabling credentials that fall out of active use. Dormancy programs practically implement this guidance, making compliance a reality.

The real advantage is once you start measuring these efforts, proving your success becomes straightforward. Focus on metrics like:

- Dormant Account Closure Rate: How many dormant accounts are cleared each month compared to the overall backlog.
- Mean Time to Remediate (MTTR): How quickly are you cleaning up dormant entitlements once identified.
- Reduction in UAR Scope: What's the percentage drop in user-access review efforts after eliminating dormant access.
- False Positive Rate: How often are you incorrectly flagging active service accounts as dormant.

Executives and auditors both appreciate charts showing clear improvement, by shrinking backlogs, quicker remediation cycles, and fewer false alarms. Demonstrating these trends visually can unlock funding and support, turning compliance from a burden into an organizational win.

How Dormant Access Cleanup Reduces UAR Fatigue

User Access Reviews (UAR) are notorious for overwhelming reviewers with spreadsheets filled with entitlements, many of which have not been used in months or even years. Dormant accounts and entitlements quietly pile up, inflating review efforts and burying critical access decisions under a mountain of irrelevant data.

Here is why clearing dormant access before running your next UAR campaign is a game-changer:

- **Massive Reduction in Review Volume:** By proactively removing the unused entitlements, you cut down the review workload significantly, making the process quicker and more focused.
- **Better Reviewer Engagement:** When reviewers are asked to approve access they rarely or never see in action, they naturally default to rubber-stamping. Eliminating dormant entitlements ensures that each item reviewed is genuinely relevant, improving the quality of decisions.
- **Faster Reviews:** Less noise means faster review cycle times. Your audit and engineering teams can quickly move through streamlined reviews, freeing valuable hours to focus on strategic tasks.
- **Reduced Audit Burden:** Auditors appreciate streamlined UAR processes where each reviewed entitlement can be easily justified. Cutting dormant access before reviews reduce auditor questions and helps you demonstrate clear governance controls.

In short, removing dormant access before UAR cycle doesn't just simplify the review, it transforms it from a compliance headache into a meaningful security control. This ultimately saves time, effort, and resources.

Practical Recommendations for Action

In summary, organizations should consider the following steps to tackle dormancy:

1. Integrate logs from Active Directory, Unix/Linux systems and cloud platforms into your SIEM for holistic monitoring.
2. Deploy modern IGA solutions with built-in analytics and automation features to continuously detect and remediate dormant accounts.
3. Regularly review policies, removing unnecessary privileges promptly upon completion of projects or departure of contractors.
4. Implement stringent lifecycle management policies for both human and non-human identities, ensuring dormant access does not accumulate unnoticed.


By proactively managing dormant accounts, enterprises must mitigate insider threats, reduce audit complexities and enhance their overall cybersecurity resilience.

Dormant accounts which are not managed and remediated will inevitably be exploited. Act before it is too late.

About the Author

Durgaprasad Balakrishnan is an independent cybersecurity researcher and Director of Cybersecurity – Identity and Access Management at a leading global fintech company. With over 16 years of experience in identity architecture, access governance, and secure automation, he has led enterprise-scale IAM transformations and contributed to multiple peer-reviewed cybersecurity initiatives. He actively participates in research communities and helps organizations design identity-centric security strategies for regulated and high-risk environments.

He can be reached via [LinkedIn](#).





Personal Cybersecurity: Benefits for Individuals and Your Business

By Paul Pioselli, Founder & CEO, Solace – Truly Personal Cybersecurity

In today's hyperconnected world, cybersecurity isn't just a corporate IT issue. It's a personal responsibility that impacts your workplace more than you may realize. When employees reuse passwords, fall for phishing scams, or use unsecured devices, it opens the door for cybercriminals to target entire organizations.

It is becoming more apparent that personal cybersecurity habits often carry over into the workplace. We're living in an era where individuals who actively protect their personal digital footprint help strengthen the security of the companies they work for. Let's explore how personal habits can impact business security, especially in the era of remote and hybrid-work.

Why Personal Cyber Gaps Put Businesses at Risk

You might think your employee's social media password or home Wi-Fi network has nothing to do with your business. But attackers don't make that distinction. They look for weak links anywhere.

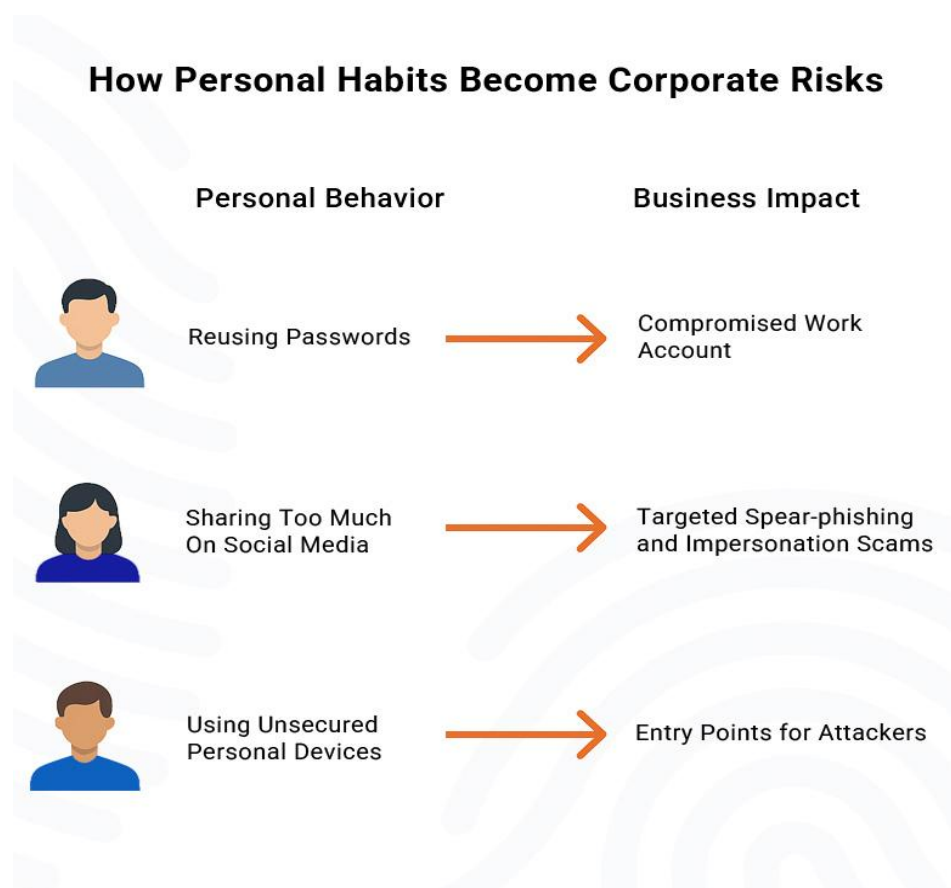
Consider this: *74% of data breaches in 2023 involved a human element, such as social engineering or credential misuse*, according to Verizon's Data Breach Investigations Report (DBIR) 2023.

Let's look at a real-world example:

An employee reuses the same password on Facebook as they do in your company's CRM system. Their Facebook password gets exposed in a breach. A hacker cross references these credentials and logs into the company CRM, stealing sensitive client data. Now the company faces reputational damage, legal costs, and lost trust, all because of a personal cybersecurity failure.

How Personal Habits Become Corporate Risks

Here are some examples of how personal cybersecurity behaviors can quietly put entire companies in danger:



These aren't far-fetched scenarios; they happen every day.

Remote Work Has Blurred the Lines

The rise of hybrid and remote work has made personal cybersecurity more important than ever. Employees use personal devices to check work email, store files in cloud accounts, and attend video calls from home networks.

Many companies report that remote employees increase their company's attack surface.

Here's why:

- Home routers are rarely secured properly.
- Personal laptops may lack antivirus software.
- Kids may use the same device to browse games or unknown websites.

Suddenly, your living room becomes a weak point in a billion-dollar company's defense.

Personal Cybersecurity: The Hidden Talent Indicator

Surprisingly, personal cybersecurity is becoming a key hiring signal for forward-thinking companies. Candidates who understand digital risks and practice safe habits tend to:

- Be more careful with sensitive data
- Report phishing attempts faster
- Respect company policies on privacy and device usage

Cyber-awareness is now a soft skill, and it's a marketable one.

Hiring security-minded people isn't just good for the IT team. It creates a culture of vigilance and trust.

Business Leaders Are Taking Notice

Many businesses are investing in employee cybersecurity training, but training alone isn't enough. According to the University of Chicago, *A one-time slideshow on phishing doesn't change lifelong habits. Research finds commonly mandated training, such as annual training methods, may be insufficient; emphasizing the need for complementary approaches to better protect organizations against phishing attacks.* (New Study Reveals Gaps in Common Types of Cybersecurity Training. 2025)

What works better is individuals investing in their personal cybersecurity, where employees take the time and make an investment to secure their own digital lives. When people protect themselves better, they also protect the company.

Here are some of the business benefits of proactive personal cybersecurity:

- Secured personal accounts lead to fewer credential stuffing attacks
- Updated home devices reduce a company's exposure to remote vulnerabilities
- Private social media settings provide a lower risk of impersonation scams

Companies Should Encourage Personal Cyber Hygiene

Here are 5 practical ways companies can support personal cybersecurity:

1. Offer Personal Cybersecurity Workshops
 - a. Incentivize employees to secure home Wi-Fi, enable 2FA, and store passwords safely.
2. Provide Identity Theft Protection
 - a. Give employees access to identity monitoring or protection services or offer reimbursement for personal cybersecurity consulting.
3. Bolster Your Cyber Culture
 - a. Normalize safe behaviors. Praise those who report phishing. Make security part of daily personal life, not just corporate compliance.
4. Expand Your Audience for Cyber Tabletop Exercises
 - a. Run mock phishing campaigns to raise awareness and reduce real-world mistakes beyond your executive team.

Empowering Individuals, Protecting Companies

Cybersecurity isn't a one-person job, but it starts with each of us at the individual and family level. When employees care about their digital safety, they act more responsibly at work. That ripple effect reduces risk across the board.

It's no longer enough to secure the corporate eco-system. You need to ensure proactive personal cybersecurity and behavior changes of the person behind the keyboard.

Final Thoughts


Personal cybersecurity is no longer just for individuals and families; personal cybersecurity is business cybersecurity.


The choices individuals make at home, what apps they download, how they store passwords, which links they click, can have massive consequences in the workplace. Companies that recognize this relationship will be more secure than those that don't.

Investing in personal cybersecurity isn't just good for people. It's a smart business move.

About the Author

Paul Pioselli is the Founder & CEO of Solace – Truly Personal Cybersecurity. He's an accomplished cybersecurity leader with over 15 years of experience, previously serving as Regional CISO in International Markets for a Fortune 15 global company before founding Solace. Solace is a boutique personal cybersecurity firm based in Connecticut, focused on protecting individuals, families, executives, and professionals with personalized, concierge-level cybersecurity. Paul can be reached online at clients@hisolace.com and at our company website <https://hisolace.com>





Shadow APIs: The Silent Backdoor Undermining Application Security

The API Economy's Blind Spot

By Sandeep Dommari, Principal Architect, Ping Identity

Introduction: The API (Application Programming Interface) Economy's Blind Spot

APIs are now what hold modern applications together. APIs now manage the majority of data exchange across enterprise environments, from cloud-native microservices to mobile apps. Unquestionably, APIs facilitate agility, quicker development cycles, and large-scale integration.

However, there is a hidden cost to this agility: the growth of **shadow APIs**, undocumented, unmonitored, or forgotten endpoints that evade **DevSecOps pipeline governance**.

Shadow APIs do not appear in inventories, vulnerability scans, or compliance audits, in contrast to a known vulnerability in a published API. Attackers love to take advantage of these invisible doors. Furthermore, the trend is only getting faster.

According to **Gartner**, most enterprise data breaches will be caused by APIs by 2025. Consider the recent high-profile breaches at **Peloton and T-Mobile**, which both involved unprotected APIs exposing sensitive customer data, if you believe this to be a theoretical risk.

The conclusion is straightforward: you cannot secure your APIs if you do not know all of them.

APIs have become the connective tissue of modern applications. From mobile apps to cloud-native microservices, APIs now handle the majority of data exchange across enterprise environments. The convenience is undeniable APIs enable agility, faster development cycles, and integration at scale.

But this agility comes with a hidden cost: the proliferation of shadow APIs, undocumented, unmonitored, or forgotten endpoints that slip outside the governance of **DevSecOps** pipelines.

Unlike a known vulnerability in a published API, shadow APIs don't show up in inventories, vulnerability scans, or compliance audits. They are invisible doors attackers love to exploit. And the trend is only accelerating.

Gartner predicts that by 2025, APIs will account for the majority of enterprise data breaches. If you think this is a theoretical risk, consider the recent high-profile breaches at T-Mobile and Peloton, both of which involved unprotected APIs exposing sensitive customer data.

The takeaway is simple: if you don't know all your APIs, you can't secure them.

Shadow APIs: What Are They?

Shadow APIs naturally arise from the rapidity of contemporary development; they are not intentionally malicious:

- **Forgotten Test/Dev Endpoints:** After an endpoint is released, it is never taken down, even though a developer spins it up for debugging. Older APIs that were never incorporated into the company's current API management layer are known as unmonitored legacy APIs.
- **Microservice Sprawl:** As serverless and container-based architectures become more common, services reveal APIs that are frequently missed in documentation.
- **Third-Party Integrations:** APIs introduced by outside vendors or SaaS products have the potential to get around enterprise controls.

Shadow APIs can be compared to a building's abandoned doors; some may be locked, but others are left ajar.

Actual Vulnerabilities: When APIs Turned Into the Attack The vector

1. T-Mobile (2023)

37 million customers' personal information, including phone numbers, account details, and billing addresses, was stolen by hackers using an exposed API. The worst part? For some queries, the API didn't even require authentication.

2. Peloton (2021)

Even when set to private, a security researcher discovered that Peloton's API revealed users' location, age, gender, and private profile information. An example of a classic shadow API error was the absence of appropriate authorization checks.

3. Parler (2021)

Attackers used a poorly secured API without rate limiting to scrape millions of posts and videos from the Parler platform.

Every case highlights the same issue: shadow APIs are more vulnerable to attacks because they are not subject to the same scrutiny and hardening as "official" production APIs.

Why Conventional Defenses Are Ineffective

Many businesses believe their API gateway or WAF (**Web Application Firewall**) offers adequate protection. However, these tools are only able to safeguard APIs that they are aware of. Due to their lack of documentation, shadow APIs avoid:

- **API inventories** - because they're undocumented.
- **Gateway security** - because developers directly expose them.
- **Frameworks for testing** - because they are not part of the CI/CD pipeline.

For this reason, shadow APIs are more than just a technical issue, they are a governance one. Conventional defenses are predicated on visibility. Invisibility is ideal for shadow APIs.

How DevSecOps Creates Shadow APIs

Software delivery is accelerated by modern DevSecOps, but uncontrolled speed encourages shadow APIs:

- **Feature Velocity**: APIs are deployed by development teams more quickly than security teams can catalog them.
- **Inadequate Documentation Culture**: OpenAPI/Swagger APIs are pushed without specifications.

- **Siloed Teams:** No API inventory is shared by developers, security, and operations.
- **Cloud-Native Complexity:** APIs are dynamically exposed by ephemeral services in Kubernetes clusters.

The irony? Modern software has blind spots because of the very techniques that make it quick and flexible.

The Developer's Guide: How to Get Rid of Shadow APIs

Shifting API security to the left and integrating visibility into the DevSecOps pipeline are necessary to handle shadow APIs. This is a useful road map:

1. API Inventory & Discovery

- To find undocumented endpoints, use automated discovery tools (such as Salt Security, Noname, Traceable, and 42Crunch) to scan traffic.
- As part of pull requests, require OpenAPI/Swagger specs for every API.
- Maintain a centralized API catalog that the development, security, and operations teams can access.

2. Include API Security in CI/CD tools for embedding linting that checks API schemas prior to merging.

- Execute API security tests in pipelines, such as checking for missing rate limiting and broken auth.
- Deployments without the appropriate API documentation are automatically rejected.

3. Visibility at Runtime

- Install API traffic analyzers at points of entry (API gateway, service mesh such as Istio, or Kubernetes Ingress).
- Signal APIs that get around gateways.
- Even for internal APIs, use rate limits, schema validation, and JWT verification.

4. Legacy API Management

- Examine current APIs and categorize them as secure, refactor, or retire.
- For APIs that cannot be modernized right away, develop compensating controls (MFA, IP restrictions).

5. Change Culture, Not Just Equipment

- Document, test, and monitor APIs just like you would production code.
- Assign ownership of the API to each product team. Encourage developers to shut down endpoints that aren't being used.

A Checklist for API Security for Developers

When releasing an API, consider the following:

- Is there a published OpenAPI specification for it?
- Does it have an API gateway integrated?
- Are rate limits and auth enforced?
- Is it observed during production?
- Do we have a plan in place to retire deprecated APIs?

You might have just created tomorrow's shadow API if any of these questions have a "no" response.

New Instruments and Methods

The industry is catching up, which is good news:

- **Service Meshes (Istio, Linkerd):** Provide telemetry and authentication at the API level without requiring code modifications.
- **AI-Driven Discovery:** Machine learning is now used by tools to map API traffic patterns and identify irregularities.
- **Continuous Validation:** API security checks are now natively integrated into Jenkins pipelines, GitLab CI, and GitHub Actions.
- **SBOM for APIs:** Similar to software bill of materials (SBOM), some businesses are requesting an API BOM in order to monitor deployments.

The Business Case: The Importance of CISOs and CTOs

Shadow APIs pose a risk at the board level in addition to being a pain for developers:

- **Regulatory Impact:** PII leaks from exposed APIs are against the CCPA, GDPR, and HIPAA.
- **Reputation Damage:** Overnight, consumer trust is undermined by a single shadow API breach.
- **Operational Cost:** Proactive governance is far less expensive than post-breach forensic investigations.

The takeaway for CISOs is unmistakable: API security is now a strategic business risk rather than a subdomain of app security.

Conclusion: From Shadow to Spotlight

Every enterprise is now an API enterprise. APIs are the currency of modern business, regardless of whether you're developing SaaS integrations, mobile apps, or microservices.

However, APIs turn into liabilities in the absence of visibility. Shadow APIs are the covert backdoor that compromises application security; they are not a specialized issue.

The way forward necessitates:

1. **A discovery-first mentality:** you cannot safeguard what you do not understand.
2. **Integration with DevSecOps:** security needs to be a part of the pipeline, not after deployment.
3. **Shared responsibility:** architects, security engineers, and developers all have a part to play.


Speed-driven APIs shouldn't be used as entry points for security flaws. We promote innovation without compromising trust by exposing shadow APIs.

About the Author

Sandeep Dommari is a Senior Cybersecurity Architect and IAM Strategist with over 18 years of experience designing secure access frameworks across Fortune 100 enterprises. His work focuses on application security, adaptive identity, and building secure-by-design architectures for critical industries.

Sandeep can be reached online at sandeep.dommari@gmail.com





The Rise of AI-Driven Credential Stuffing: Why Identity and Access Management (IAM) Alone Can't Save You

When Bots Learn to Think Like Humans

By Sandeep Dommari, Principal Architect, Ping Identity

Introduction: When Bots Learn to Think Like Humans

Security teams dismissed credential stuffing as "noisy bot traffic" in 2012. It has now developed into one of the world's most lucrative, scalable, and AI-powered threats.

Consider the recent spate of hacks at Nintendo, Zoom, and Spotify. Attackers only needed to use usernames and passwords that had been stolen from unrelated leaks and allow bots to test them across millions of accounts; they didn't even need to breach the companies directly.

Attackers no longer simply spray stolen credentials thanks to generative AI. They are so good at imitating human behavior that they are fooling even sophisticated Identity and Access Management (IAM) systems and conventional bot detection tools.

The Evolution of Credential Stuffing Beyond "Just Bots"

In the past, credential stuffing involved using speed bots to brute force login forms with compromised usernames and passwords.

- They were frequently prevented by security measures like IP blacklisting, velocity checks, and CAPTCHA.
- The New World with AI: AI-powered bots dynamically modify attack velocity, device fingerprints, and session patterns through reinforcement learning.
- Attackers repurpose tools such as open-source machine learning libraries to teach bots to "look" like actual users.
- Credential stuffing campaigns now circumvent MFA by taking advantage of push notification fatigue or weak SMS-based factors (also known as "MFA bombing").

Real-world illustration:

An AI-driven attack against a multinational retail behemoth in 2023 involved bots that mimicked human shopping behavior by rotating IP addresses across mobile networks, simulating device orientation sensors, and even varying login attempts over several weeks. There was nothing unusual in the IAM logs. The fraud bill? Chargebacks and account takeovers totaling more than \$20 million.

The Business Impact: The Importance for CEOs and Boards

Credential stuffing is more than just an IT annoyance to a CEO or board member. There are repercussions from this business continuity risk:

- **Brand Damage:** Even if credentials originated from another breach, the victim always holds the company accountable when a user account is compromised.
- **Fraud Costs:** Mostly as a result of credential stuffing, airlines report \$750 million in loyalty program fraud each year.
- **Operational Disruption:** Calls for password resets and fraud investigations overwhelm customer support teams.
- **Regulatory Penalties:** In accordance with the GDPR/CCPA, there may be fines for neglecting to protect customer accounts, even from credentials that are reused.

For instance, Nintendo acknowledged in 2020 that a credential stuffing campaign had compromised 160,000 accounts. Parents and regulators were outraged when the attackers targeted children's accounts for stored credit card data.

Why You Can't Be Saved by Me Alone

IAM systems are essential, but they are not made to handle this issue.

- Although attackers can use AI to evade MFA prompts, IAM can enforce strong authentication.
- Although IAM can centralize identities, APIs and shadow apps introduce vulnerabilities.
- Although AI-driven bots produce noise that appears to be human logs, IAM offers audit logs.

IAM is reactive, which is a painful reality. Layered defenses and proactive detection are necessary for credential stuffing.

How AI Enhances Credential Stuffing

Attackers use AI as a weapon in the following ways:

1. **Behavioral Mimicry:** By recording actual user sessions, bots teach machine learning models to imitate mouse motions, geolocation switching, and typing rhythm.
2. **Learning that Adapts:** Bots use residential proxies to route if IPs are blocked; they learn from unsuccessful login attempts. They target distinct accounts if MFA is activated.
3. **Changes to Passwords:** Passwords that have been stolen are transformed into dozens of different variations by generative models ("Summer2023!" → "\$umm3r2023!!").
4. **Getting around CAPTCHAs:** The majority of CAPTCHAs are broken at scale by computer vision and LLM-powered solvers.
5. **Abuse of APIs:** Bots circumvent web defenses by directly exploiting login APIs.

Case Study: In order to evade fraud detection, a financial services company found that attackers were utilizing reinforcement learning bots that modified login attempts in real time. Before the attack was lessened, it took six months and a new bot defense solution.

Developing a Defense Outside of IAM

1. Feeds of Credential Intelligence

- Incorporate threat intelligence that keeps an eye on dark web dumps and sends out alerts when user credentials show up.
- For instance, businesses that use these feeds proactively reset exposed accounts following the LinkedIn leak.

2. AI-Powered Bot Detection

- Use anomaly detection at the edge (WAF/CDN), which examines behavioral patterns such as device fingerprinting, velocity, and mouse movement entropy.

- Best practice: Integrate in-house tuning with vendor solutions (like PerimeterX and Akamai Bot Manager).

3. Authentication that Adapts

- Get rid of static MFA. Make use of risk-based policies:
- Only use step-up authentication in cases where the session risk score is high.
- Take into account device history, IP reputation, geolocation, and impractical travel.

4. Security Controls for APIs

- Keep web login flows and login APIs separate.
- As an illustration, apply rate limits and schema validation, especially for mobile app login APIs.

5. Transparency & User Education

- Inform clients about the dangers of password reuse and send out password breach notifications when credentials are found to be reused.
- Openness fosters trust.

A Guide for CEOs and CISOs

Phase 1: Visibility & Discovery

- Map every login endpoint, including mobile flows and APIs.
- Benchmark failed login ratios: suspicious increases frequently indicate credential stuffing.

Phase 2: Incorporate Controls

- Enhance the current IAM with threat intelligence, bot detection, and adaptive authentication.
- Enhance IAM rather than completely replace it.

Phase 3: Ongoing Examination

- Conduct red-team drills that replicate credential stuffing in particular.
- Apply the concepts of chaos engineering to test the resilience of your login flow in the event of an attack.

Phase 4: Involvement with Businesses

- Report credential stuffing as revenue loss and customer fraud rather than a "login failure issue."
- Frame board talks about revenue and reputational risk.

Conclusion: The Arms Race in AI Has Started

Automation versus defense has always been at the center of credential stuffing. However, AI has ushered in a new era of intelligent, adaptable, and outsmarting robots.

The lesson for CISOs and CEOs is straightforward: IAM is important, but not enough.

Those who invest in layered defenses, test resilience frequently, and treat credential stuffing as a systemic business risk will emerge victorious.

If not, attackers will be able to walk through your digital front door rather than just knock on it.

About the Author

Sandeep Dommari is a Senior Cybersecurity Architect and IAM Strategist with over 18 years of experience designing secure access frameworks across Fortune 100 enterprises. His work focuses on application security, adaptive identity, and building secure-by-design architectures for critical industries.

Sandeep can be reached online at sandeep.dommari@gmail.com





The Future of Privileged Access is Vault-Free

By Yaron Kassner, Co-Founder & CTO, Silverfort

The cybersecurity world was jolted by the recent announcement that Palo Alto Networks will acquire CyberArk in a landmark deal valued at approximately \$25 billion. Beyond the financial scale of the transaction, this acquisition marks a shift in how the industry views identity security. The recent acquisition validates what we've been emphasizing already: **identity is both the first and last line of defense**, and it demands its own dedicated security layer. Just as we've seen in other domains like endpoint and cloud security, protecting identities requires an end-to-end **platform**—one that offers unified visibility, intelligent insights, and inline protection.

CyberArk, rooted in Privileged Access Management (PAM), has expanded its identity capabilities in response to market needs. But this acquisition surfaces a more critical and timely question:

What is the future of PAM?

We believe we're witnessing the beginning of a transformation. A future where securing privileged access will no longer revolve around a **vault**. In fact, vault-based approaches will no longer be the *primary* method of enforcing privileged access security.

This shift parallels transformations we've already seen in other areas of cybersecurity:

- **Cloud security** has gone agentless, replacing intrusive deployments with lightweight, API-based visibility.
- **Multi-Factor Authentication (MFA)** does not require code changes or proxies as in the past. Today it is enforced by the identity provider or an extension to the identity provider.
- **Identity protection** platforms now enforce policies in real time without injecting keys or passwords, reducing the attack surface.
- **Network security** moved away from physical firewalls and VPNs to Zero Trust Network Access (ZTNA), which grants access dynamically based on identity, context, and posture.

In each of these cases, the core idea was the same: move away from securing **secrets or infrastructure**, and instead focus on securing the **access itself**. PAM is now undergoing a similar evolution.

The Problem with Vault-Based PAM

Vaults were introduced as a way to protect the credentials used by privileged accounts—admin usernames and passwords for servers, databases, switches, and more. The premise was sound: don't let users know or reuse powerful passwords. Instead, let them retrieve credentials from a secure vault when needed, and rotate those passwords after use.

But in practice, vault-based PAM creates several problems:

1. **It secures the credential, not the access.** Once a user retrieves the credential, the vault's protections end. That password can be stolen from memory, logged by malware, misused by insiders, or intercepted in a man-in-the-middle attack. The access itself isn't protected—just the storage of the password.
2. **It's operationally complex.** Vault-based PAM introduces major friction into workflows. Changing how users log into systems—redirecting them through a proxy, forcing them to check out passwords, re-authenticate constantly—often requires training, workarounds, or exceptions. On the NHI front, to rotate service account credentials multiple approvals are typically required and careful work to avoid breaking changes. This change in behavior complicates adoption and makes PAM deployments **time-consuming and expensive**. Many organizations take **years** to roll out PAM at scale, especially in hybrid environments where legacy systems, service accounts, and third-party access all require separate configurations.
3. **It's not breach-proof.** Vaults themselves are high-value targets. Attackers know that compromising a vault can yield credentials for the most sensitive systems in the organization. We've seen real-world breaches that prove this. In a **high profile** 2022 breach, the attacker reportedly gained access to the company's privileged access vault by harvesting credentials and tricking an employee into approving MFA requests. Once inside, the attacker had access to admin tools, infrastructure, and sensitive data. In other incidents, attackers have exploited vault misconfigurations, API tokens, or integration weaknesses to escalate their access. The idea that vaults are unbreachable is no longer tenable.

4. **It creates a false sense of security.** Security teams often assume that rotating credentials and limiting access to the vault is enough. But if the password is still being handed to the user—even for a short time—it can still be exfiltrated or abused. The security controls (like MFA, session recording, or approval workflows) are tied to the vault, not to the privileged access itself. Once the login is done, there is no additional enforcement point to apply security controls.

Vault-centric PAM worked well in the era of static infrastructure and long-lived accounts. But today's IT environments are dynamic, distributed, and identity-driven. Simply protecting credentials in a vault is no longer enough.

From privileged account management to privileged access security

The real opportunity—and what defines the vault-free future—is to shift from **managing privileged accounts** to **securing privileged access**.

In this model, organizations no longer rely on permanent accounts with vaulted passwords. Instead, privileges are granted dynamically, just-in-time, and removed as soon as they're no longer needed. Access is brokered and monitored in real time based on user identity, context (device, location, time), and policy.

This eliminates many of the risks associated with vault-based PAM:

- There is no standing credential to steal or reuse.
- The change to user behavior is minimal - no login disruption, and no password checkout process.
- All access is tightly monitored and tied to a verified identity.
- Even if the attacker gains hold of the password, the access is still secured and the attack can be stopped there.

This model also extends seamlessly to **non-human identities** (NHIs)—like service accounts, scripts, AI agents and automation tools—which now make up the majority of privileged access in most organizations. Rather than managing thousands of long-lived credentials for these entities, organizations can enforce policies that allow specific systems to initiate privileged access under strict controls, without static secrets. As NHIs become more manageable through identity providers, cloud-native tools, and runtime enforcement, the vault-free approach becomes both more feasible and more secure.

Identity-Centric Access: A More Secure Approach

This shift toward privileged access security is made possible by technological advances in identity security. Organizations can now apply strong security controls at the identity layer—enforcing **MFA, risk-based policies, session monitoring, and just-in-time elevation**—*without* injecting credentials or modifying infrastructure.

In fact, modern platforms can secure privileged access in a way that's:

- **Proxyless** – doesn't require routing all of the network traffic through a gateway or rewriting apps.
- **Credential-free** – avoids injecting or exposing privileged credentials.
- **Inline & real-time** – dynamically responds to access attempts with adaptive policy decisions.

This architectural shift allows organizations to apply Zero Trust principles to privileged access—validating every request continuously, applying least privilege policies, and responding to anomalies instantly.

And it aligns with how security teams want to work: reducing the attack surface, minimizing user disruption, and simplifying operations.

Will Vaults Disappear?

Vaults will remain part of the privileged access landscape for the foreseeable future. Some systems will continue to require passwords. Some compliance requirements will mandate secure storage of credentials. And in certain break-glass or legacy scenarios, having a vault as a fallback mechanism still makes sense.

But vaults will no longer be the **primary** way organizations secure privileged access. Instead, the center of gravity will shift to **real-time, identity-aware controls**—a model that doesn't rely on handing users credentials, and doesn't require those credentials to exist in the first place.

We're already seeing this transition unfold. Modern identity security platforms are being used to enforce granular access controls for privileged sessions across cloud and on-prem environments. These controls—based on who the user is, what resource they're accessing, and under what context—are more precise, more scalable, and more secure than vault-based approaches.

And importantly, they're **faster to deploy and easier to manage**, because they don't require users to change how they log in or IT teams to redesign their environments.

Looking Ahead

The future of privileged access is vault-free. Vaults served a critical function in an earlier era. But as identity becomes the new perimeter, and access becomes the control point, it's time to move on.

Security leaders who want to reduce risk, accelerate zero trust adoption, and simplify their operational burden should begin by asking: *Do I need to protect this password, or can I eliminate it altogether?*

By shifting the focus from accounts to access, we can finally secure identities in a way that's invisible to users, resistant to breaches, and built for the dynamic environments of today—and tomorrow.

About the Author

Yaron Kassner is Co-founder and CTO of Silverfort, where he has worked for eight years. Previously, he was a data science consultant for Cisco and has also worked as a software development engineer at Microsoft. He holds a PhD in computer science from the Israel Institute of Technology.

Yaron can be reached online on [LinkedIn](#) and at our company website <https://www.silverfort.com/>





How In-Memory Vulnerability Scanning Boosts Enterprise Linux Security

By Eric Hendricks, Product Manager for the Radar Vulnerability Scanner at TuxCare

Enterprise Linux security faces a dangerous new reality as AI-powered threats continue to become more adaptive and aggressive toward traditional methods of defense. With the cyberthreat landscape continuing to evolve, routine vulnerability scanners and legacy systems are struggling to keep up because they're unable to differentiate between a low-risk concern or a severe threat needing immediate attention from the security team. The noise can often become overwhelming. As a result, systems get cluttered with false positives, forcing security teams to spend valuable time manually addressing alerts.

False Alarms Are Draining Time and Resources

Not only is that manual process highly disruptive, it's ultimately dangerous. While security teams and sysadmins are busy sorting false alarms, critical vulnerabilities are getting put on the backburner or missed entirely. This combination creates a wealth of opportunity for money hungry hackers looking to exploit vulnerabilities and capitalize on overwhelmed, ineffective defenses.

Relying on a reactive incident response plan can financially devastate a company trying to fight AI-generated attacks with traditional tools, which is why automating your vulnerability scanner is critical for managing future risk. Routine scanners can easily identify vulnerabilities based on installed packages. However, they often ignore what's going on in-memory: creating a major blindspot for IT teams. By only monitoring installed packages and not the live state of the system, false alarms will continue to clutter scanners.

Don't Let Traditional Vulnerability Scanners Jeopardize Operations

Sysadmins need clarity, not clutter. Reducing excessive false positives starts with implementing a vulnerability assessment solution that has the built-in intelligence for in-memory patch awareness. Not only should it be able to recognize patched systems, it should also be able to understand whether updates were applied through rebootless patching, extended support channels, or manual fixes. Operating with this level of clarity allows the scanner to rapidly triage the most urgent threats, completely eliminating the need to manually vet flagged CVE's for vulnerabilities.

With in-memory awareness monitoring comes an emphasis on rebootless patching that can streamline the identification, testing and deployment of patches across the IT infrastructure. Traditional patching processes have long required extended downtime and routine maintenance windows that pose a risk to day-to-day operations. But opting to automate the process means critical security patches can be applied to bugs and vulnerabilities directly to the kernel in real time. Not only does this help companies avoid having to pay for costly, vendor-specific live patching solutions or support packages, but it also helps satisfy a number of compliance requirements including CIS controls, NIST CSF, and PCI DSS.

AI-assisted risk analysis also enables security teams to remediate threats more efficiently while streamlining compliance reporting. With a lightweight CLI tool that allows unprivileged scans, automated scheduling and integration into existing workflows, the burden of remediation gets greatly reduced without compromising visibility or control.

Cut Through the Noise

While it cannot be expected to replace human judgement and oversight, a vulnerability scanner enhanced by AI undoubtedly serves as an essential tool for sysadmins to have in their back pocket in the fight against today's threats. From reducing risk exposure, to streamlined reporting, proactive vulnerability scanning can not only save time and resources, but also provide security teams with the freedom they need to focus on their core business.

About the Author

Eric Hendricks currently serves as product manager for the Radar vulnerability scanner at TuxCare.

Eric can be reached online at <https://www.linkedin.com/in/itguyeric/> and at our company website www.tuxcare.com





The Vibe Coding Conundrum: Development vs. Security

By Amir Kazemi, Director, Product Marketing, Cyscale

"Trust the vibes" has long been a phrase amongst new-age philosophers and manifestation experts, but should it have a place in developing code?

Vibe coding has rapidly emerged as a captivating topic among developers, igniting widespread excitement, curiosity, and fervent debate. This innovative paradigm allows developers to craft code using intuitive AI-driven tools, primarily through natural language prompts, rather than the painstaking, line-by-line syntax of traditional programming. The result is a remarkably fluid and effortless creative process, akin to achieving a state of "flow" or being "in the zone."

As developers chase speed and creativity, it's easy to get swept up in the excitement of vibe coding and its vast potential for outcomes. However, as the pace of development accelerates, a critical question arises: With vibe coding, are we just moving faster and overlooking security?

Speed Without Safeguards is a Ticking Time Bomb

Vibe coding is exposing existing structural weaknesses in how organizations handle software security. The capacity to manage code effectively has become increasingly strained, with the ratio of security resources to developers shifting from approximately 1:100 to now closer to 1:1000.

In an already fragile software ecosystem, vibe coding outstrips the safety nets meant to help protect developers who are simply not trained to detect the nuanced, hidden threats that security teams can identify. While the tools used today for vibe coding may perform some basic security testing, they're still very far from comprehensive or accurate, exposing organizations to significant risk. As code output explodes, lean security teams are left searching for a needle in a haystack that's growing by the page. The resulting skills gap and power struggle to thoroughly review changes, while developers rush to complete software, creates chaos. Flaws often slip through the cracks simply because no one has the time or eyes to catch them.

As vibe coding takes off, this is only going to get worse. But there is a fix. Companies must embed security into their everyday development, not just bolt it on at the end. This works twofold by empowering developers with practical training and tools to identify issues early and foster collaboration with the security team. Viewing security as a shared responsibility, teams can harness the power of vibe coding and stay safe. Without it, the trend could come at a devastating cost.

Security by Design: The Missing Layer in Vibe Coding

Today's apps are complex, and security risks can hide deep within this complexity. Examples of potential issues include SQL injections, cross-site scripting, leaked secrets, and supply chain vulnerabilities. When developers use vibe coding and move quickly, these risks can escalate. That's because AI-powered code can sometimes introduce insecure patterns, especially when developers lack the training to spot and fix these issues.

To prevent vibe coding from introducing these dangers, security has to shift even farther left. This means there must be continuous, automated security checks at every step, from the initial AI-generated code snippets to the final deployment and beyond. With real-time visibility and context, teams can identify vulnerabilities before they become embedded in fast-moving codebases.

The Role of Context-Aware Security

Traditional, reactive security checks won't cut it for vibe coding. Security solutions shouldn't just scan code after the fact; they must understand how the application is built to flag vulnerabilities in real time as the software takes shape.

Instant, context-rich feedback and actionable guidance from security tools can empower developers to resolve security issues *before* they result in costly downtime and financial losses for organizations. Providing developers with precise, actionable feedback within environments such as IDEs, code review

platforms, CI/CD pipelines, and other tools is the final step in making security enforcement a continuous, integrated function.


Preparing for the Future


Ultimately, the debate isn't about whether vibe coding is good or bad. And it's not about adding personnel to bridge the growing divide between developers and application security professionals. The real question is how we can adapt our tools, habits, and culture to make sure that tomorrow's software is both groundbreaking and secure without sacrificing robust protection. The broader sentiment around AI's benefits in software development continues to evolve, but let's be clear: embracing innovation means proactively adapting our security approach.

To mitigate these risks, organizations should focus on education, better tools, and building a culture where security is everyone's responsibility. Leveraging a native AI application security platform that provides features like AI Code Security Assistants can further enhance these efforts. By making security an integral part of the development process, companies can fully embrace the advantages of vibe coding while protecting their applications from new and evolving threats.

About the Author

Amir Kazemi is the Director of Product Marketing at Cyscale. Amir has an extensive background in cybersecurity and tech marketing and specializes in implementing cutting-edge AI solutions to strengthen software security and foster collaboration between development and security teams. He currently drives the evolution of Cyscale's AI-Native Application Security platform. Amir can be reached online on [LinkedIn](#) and at our company website <https://cyscale.com/>





Creating a Full-spectrum Cyber Workforce for CJADC2

By Magdalena LoGrande, Cybersecurity Engineering Fellow, Sigma Defense

The Defense Department's evolving cyber workforce framework, as outlined by the [8140 Program](#), offers a perfect opportunity to reflect on its momentous alignment with CJADC2 objectives and to highlight the foundational role of full-spectrum, cross-functional expertise to achieve its vision.

CJADC2 is an ecosystem of capabilities. It aims to connect military assets across all domains to enhance decision-making and improve operational effectiveness, emphasizing interoperability and collaboration with allies and partners. As a warfighting domain alongside land, sea, air and space, cyberspace needs a robust and integrated cyber workforce.

This was underscored by Commander of Cyber Defense Command and DISA Director Lt. Gen. Paul Stanton when [he stated](#) at the DoD Cyber Workforce Summit in April that “talent is our asymmetric advantage.”

DoD 8140 codifies the cyber workforce as an ecosystem of expertise. It supports CJADC2 by outlining a taxonomy of talent that deliberately encompasses “the entire cyberspace workforce, to include personnel who build, secure, operate, defend, and protect DoD and U.S. cyberspace resources, conduct related intelligence activities, enable future operations, and project power in and through cyberspace.”

The elements of that taxonomy are Cyberspace IT, Cybersecurity, Cyberspace Effects, Intelligence (Cyberspace), Cyberspace Enablers, Software Engineering, Data/AI.

Both CJADC2 and 8140 have a “total” concept – the former expands and unifies full-spectrum operations while the latter expands and unifies the expertise that makes that possible. This is a significant step in formalizing a broader and more dynamic view of the cyber workforce that extends beyond information assurance and cybersecurity.

It’s instructive to look at some key examples of how cyber workforce elements mesh with each other to enable CJADC2 functions:

- **Cyber as a warfighting function:** The inclusion of the *Cyberspace Effects* element in the framework is the explicit acknowledgement of cyber as a defensive and offensive weapon, where operators not only protect and preserve communication networks, information systems, and digital enclaves, but also execute force projection activities as directed. The *Intelligence (Cyberspace)* element is a DCO/OCO enabler, as planners and analysts rely on the integration of sensing and cyber, especially at the tactical edge and in highly contested environments.
- **The resourcing matrix:** Integration of command and control does not only mean integration of systems and networks, where *Cyberspace IT* plays a foundational engineering implementation and sustainment role, but also stakeholder integration at the executive leadership level, where *Cyber Enablers* (capability sponsors, portfolio managers, legal advisors), align acquisition and development strategies and adopt novel approaches to tackle the hard problems posed by disjointed data architectures and perceived policy barriers, leveraging the best that agile commercial innovation and open standard models (MOSA/SOSA) have to offer.
- **Purposeful data curation for information advantage** resides at the center of the *Data/AI* interplay to prevent Garbage In/Garbage Out effect, fine tune LLMs to mission sets, and enable reuse and interoperability and making disparate data flows “make sense” to the *Intelligence (Cyberspace)* elements (both producers and consumers) across all domains via AI-enabled processes, which are in turn built upon *Cyberspace IT* and *Software Engineering* layers.
- **Data accessibility:** *Software Engineering* acts as the force multiplier that threads together modules and segments across the E2E data logistic chains (as software-defined functions support dynamic configurability in near real time); unleashes data fabric interoperability (e.g., APIs); and enables bespoke data visualization (e.g., UI/UX). Embedded in Software Engineering is the DevSecOps practice and its ability to quickly deliver mission relevant code to the warfighter fast and frequently.
- **Cybersecurity for the Data Logistics Chain:** Since CJADC2 is founded on the concept of a multi-network, multi-lateral, multi-contingent ecosystem, the *Cybersecurity* element will drive a Zero Trust (ZT) design that intelligently reduces the threat footprint for critical systems and grants network segment data access to only those entities identified as needing said access utilizing policy enforcement points across the chain. That also requires not just the presence of a ZT-trained *Cybersecurity* workforce but also significant active stakeholder involvement (going back to our executive *Cyber Enablers*) in the development and prioritization of ZT-focused solutions that do not create new stovepipes. Also critical for CJADC2, as a multi-classification architecture,

are purposeful, well-thought-out deployments of cross-domain solutions when and where they are needed.

An operational ecosystem demands a workforce ecosystem where all cyber elements are present, strong and actively engaged with each other, eliminating the proverbial silos that have been a deterrent to the CJADC2 approach. Considering human capital through a total cyber workforce is imperative — that total concept should remain at the forefront of our minds as CJADC2 implementations continue to realize and mature beyond the concept phase.

About the Author

Magdalena LoGrande currently serves as the Cybersecurity Engineering Fellow at Sigma Defense Systems, LLC, where she provides subject matter expertise to engineering teams, helping them create actionable design and implementation plans that incorporate cybersecurity as a critical component. She has supported DoD and IC customers in various roles across the full spectrum of systems security engineering and Risk Management Framework (RMF) activities since 2012. This experience has given her an opportunity to witness and harness the power of cybersecurity as an embedded engineering practice. As a thought leader in the field, she has also served on the AFCEA International Cyber Committee since 2022, and co-authored white papers, such as [A Cyber Curriculum for a Digital Workforce](#) and [Secure by Design – Next Steps](#).



Magdalena can be reached online at magdalena.logrande@sigmadefense.com and at our company website <https://sigmadefense.com/>



Cyber Risk Management for FinTechs in the 21st Century

Adapting to AI, Cloud, and IoT Disruption

By Oreoluwa Joda, Cyber Risk Manager, Fiserv

As Fintechs accelerate their adoption and integration of emerging technologies like Artificial Intelligence (AI), Cloud computing, and the Internet of Things (IoT), they are challenged by an increasingly complex and volatile cyber threat landscape. These technologies, while transformative, also introduce new and often vaguely understood risks, including but not limited to algorithmic manipulation, exposure of data in dynamic cloud environments, and widespread vulnerabilities in connected devices. Traditional cyber risk management frameworks no longer suffice in addressing the speed, scale, and interconnectivity of these innovations. Thus, Fintech companies must implement comprehensive and advanced risk management strategies to protect their data and digital infrastructures.

This article highlights the need for adaptive, technology-driven security approaches that encompass traditional methods. It explores how AI and other key emerging technologies are redefining the fintech industry, how cloud systems require continuous monitoring, and how IoT expansion increases vulnerability at endpoints. It also explores ways in which Cyber risk management can be enhanced in the fintech industry, via transition to proactive and intelligent defenses to protect the digital ecosystems effectively.

The Next Wave: AI and Other Technologies Transforming the FinTech Industry



In a survey conducted by Tribe Payments, results show that about 67% of fintechs believe that AI will have the most significant impact on the sector in the next five years. **APIs, Blockchain technology, AI, Low-Code/No-Code development, and Edge Cloud Computing are some key technologies redefining the 21st Century FinTech industry.**

Notably, about 90% of fintechs use Application Programming Interfaces (APIs), making it the most widely used emerging technology in the industry, fueled by the rise in Open Banking.

About 20% of fintechs have integrated Blockchain technology, 70% of fintechs already use AI, 16% of these organizations use Low code to accelerate the pace of code development, and 10% have adopted edge computing (Tribe Payments, 2025).

Fintech APIs facilitate the communication between financial systems, driving online payments, bank integrations, and digital wallet transfers. In FinTech, APIs serve as a connection between front-end applications, like mobile banking apps, and back-end financial systems such as payment processors. Software developers provide these APIs to enable financial institutions, startups, and developers to build innovative financial products.

Blockchain improves the efficiency and transparency of peer-to-peer payments. It also addresses security issues and minimizes fraud in payment systems, making transactions safer and more reliable. Blockchain makes auditing easier by providing a transparent, immutable record of transactions, simplifying the tracking and verification of financial data. It automates record-keeping, lowering costs and enhancing accuracy. Moreover, through cryptographic techniques and decentralized consensus, blockchain ensures secure identity verification, minimizing fraud and increasing trust in digital transactions.

Artificial Intelligence in the fintech sector is used for threat detection and prevention, the automation of Incident Response (IR), fraud detection, risk assessment and vulnerability management, as well as behavioral analysis of insider or external cyber threats.

Low-Code and No-code Development are fast rising in the contemporary Fintech industry, resolving software development challenges. Low-code development platforms allow 61% of users to complete custom applications on schedule, within budget, and according to requirements.

With features like drag-and-drop tools, pre-installed templates, and simplified interfaces, low-code platforms have enhanced software development.



Edge Cloud Computing allows for data processing at the “edge” of the network, via the device or a local server, reducing latency (Microsoft, 2025). In the fintech industry, Security and regulatory compliance laws like GDPR and PCI-DSS enforcing strict data protection must be adhered to, and non-compliance can result in significant fines and reputational harm.

Edge computing aids compliance by processing data locally and minimizing the need to transmit sensitive information over networks. This approach enhances data security, supports compliance with data residency rules, and lowers the risk of cyber threats.

The Realm of the Cloud in Fintech Operations

The transition to cloud-native environments has transformed how fintech companies operate by enabling scalability, agility, and cost efficiency. However, this transformation has also increased the potential for

cyber-attacks. Common and avoidable causes of cloud-related breaches include misconfigurations, unsecured data repositories, and poorly managed user access controls.

Cloud infrastructure often involves multiple service providers and complex deployment models, such as hybrid or multi-cloud architectures. Each additional layer of abstraction and third-party involvement makes it more challenging to maintain consistent security controls. Without full visibility into these environments, it is difficult to assess real-time exposure, enforce compliance, and respond quickly to incidents.

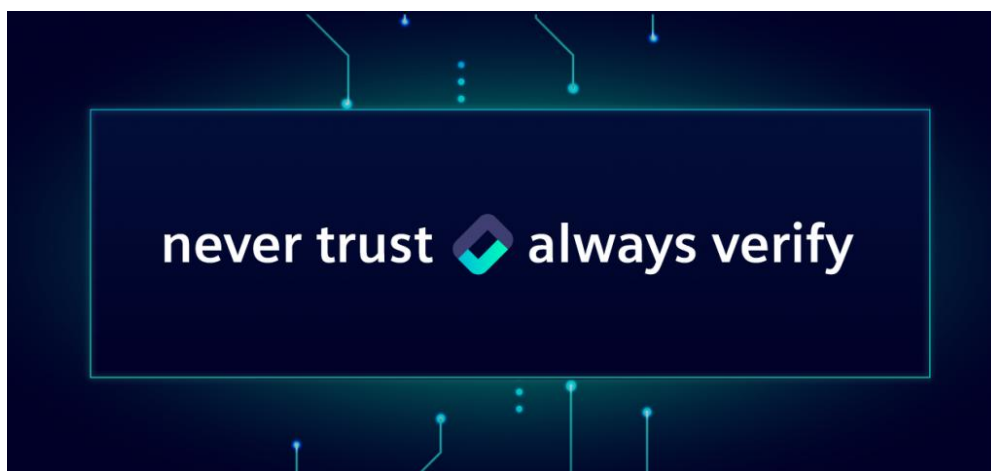
Zero Trust, Full Protection: Monitoring Encryption & Data Protection in the Cloud

The Zero Trust Cybersecurity framework is established on the principle that no user, device, or system should be trusted by default- even within an organization's internal network.

Fintechs transmit, store, and process sensitive data in the Cloud Environment including but not limited to- Personally Identifiable Information (PII) and Payment Card Information (PCI) based on their product and service offerings.

With Zero Trust, stringent verification must occur before access to any DB or application in the Cloud, encryption protocols must be continuously monitored for Data-At-Rest (DAR) and Data-In-Transit is required.

Understanding How Zero Trust Architecture Functions



The Zero Trust model is founded upon the National Institute of Standards and Technology (NIST) 800-207 framework established on three fundamental principles- "Never Trust, always verify", Least Privilege Access, and Assume Breach.

The implementation of a Zero Trust framework involves incorporating advanced technologies, including; Risk-based Multi-Factor Authentication (MFA)- verifies user identities, Identity Protection- for user

identities and systems security, Advanced Endpoint Security- protects endpoints and prevents unauthorized access, and Cloud workload technology- protects workloads across cloud platforms, reducing the risk of security breaches.

Through real-time visibility, automated alerts, and ongoing compliance monitoring in the cloud environment, Fintechs can detect security weaknesses early, prevent potential breaches, and retain complete oversight of their cloud assets, ensuring strong, end-to-end protection in the rapidly evolving digital landscape.

Addressing IoT-Driven Cyber Risks in FinTech Organizations



The widespread use of IoT devices within fintech infrastructure, ranging from payment terminals and biometric scanners to smart office systems, has not merely enhanced innovation, but also heightened risk.

The interconnected nature of IoT ecosystems implies that, once compromised, a single compromised device can serve as an open door for lateral movement across an organization's entire digital terrain. This is more concerning in fintech operations given the collection, usage, processing and storage of sensitive financial data, storage of client records, and interconnection with authentication mechanisms.

Visibility, segmentation, and lifecycle control are some significant components of an effective risk management strategy in IoT contexts. Organizations must maintain a comprehensive and accurate inventory of all connected devices, consistently monitor network traffic for anomalies, and enforce stringent access controls based on device/ application functions. Additionally, vendor agreements and partnerships should include clear guidelines regarding firmware support, vulnerability patching, and end-of-life protocols.

Bridging Technology and Strategy in Cyber Risk Governance

Modern fintech leaders should regard cybersecurity as a business strategy and tool for growth, not a back-end function. This requires transitioning from a reactive security posture to a proactive, strategic risk governance approach. Cyber risk must be viewed through the lens of enterprise risk, with direct connection to business continuity, customer trust, and regulatory compliance.

A successful cyber risk governance model integrates operational security, legal obligations, and board-level oversight. This model also fosters collaboration between IT, legal, compliance, and executive teams, harnessing a shared understanding of risk tolerance, response plans, and business impact.

AI for Effective Cyber Risk Management

By leveraging AI-driven tools, organizations can identify, prevent, and address cyber threats before they escalate into significant security breaches.

As cybercriminals adopt more sophisticated techniques like ransomware, phishing, and zero-day exploits to bypass security defenses, the need for a more intelligent and proactive cybersecurity strategy becomes more important. AI-based cybersecurity systems can rapidly process large amounts of data, identifying indicators of potential cyber threats.

Unlike traditional security solutions that follow fixed rules, AI consistently learns from incoming data, improving its ability to anticipate and prevent attacks. The integration of machine learning and automation also allows organizations to detect threats in real time, respond more promptly to incidents, and reduce the workload for security teams. AI is essential for strengthening cyber defenses by minimizing human mistakes and increasing the precision of threat detection.

Limitations of AI

While AI holds significant potential in Cyber Risk Management, it is not without its limitations. Relying solely on automated systems can result in overlooked vulnerabilities, especially when these systems encounter novel or highly targeted threats outside their training parameters. In such cases, the algorithms may fail to accurately assess risk, flag anomalies, or understand the full scope of an evolving threat landscape.

Another significant concern is the quality and diversity of data used to train AI systems. If the foundational data is incomplete, not current, or biased, the outcomes produced by AI tools may reflect these shortcomings, leading to false positives, unidentified vulnerabilities, or misdirected security responses.

In high-stakes financial environments, where precision is critical, such errors can result in significant operational or reputational damage.

Conclusion

In conclusion, Cyber Risk Management has become a strategic imperative for FinTech organizations navigating the complexities of AI integration, Cloud computing, and IoT expansion. These technologies, while enabling innovation and operational efficiency, simultaneously introduce unprecedented vulnerabilities and regulatory challenges. To remain resilient and competitive, FinTechs must adopt adaptive, risk-informed frameworks that emphasize real-time monitoring, data governance, Zero Trust Architecture, and continuous compliance. As digital ecosystems evolve, the security posture of FinTechs must stay ahead- balancing agility with accountability to protect sensitive financial data, preserve consumer trust, and ensure long-term sustainability in an increasingly interconnected world.

References

- Cardona, M. (2021). *How to Secure IoT in Financial Services? - Perspectives*. Perspectives. https://www.paloaltonetworks.com/perspectives/how-to-secure-iot-in-financial-services/?utm_source
- Edge Computing | Microsoft Azure. (2025). Microsoft.com. <https://azure.microsoft.com/en-us/resources/cloud-computing-dictionary/what-is-edge-computing#>
- Fintech 5x5. (2025). Tribepayments.com. <https://www.tribepayments.com/tribe-blog/67-of-fintechs-believe-that-ai-will-have-the-biggest-impact-on-sector-in-the-next-five-years>
- Harnessing Artificial Intelligence (AI) For Enhanced Cyber Risk Management - Brandefense*. (2025). Brandefense. <https://brandefense.io/blog/drps/ai-for-cyber-risk-management/>
- Krissansen, J. (2025). *What are FinTech APIs? Definition and how they work*. Bill.com; BILL. <https://www.bill.com/learning/fintech-api>
- Low-Code Adoption Statistics 2025: Trends, Forecasts, Insights*. (2025). Akveo.com. <https://www.akveo.com/low-code-adoption-statistics>
- Perez, M. (2025). *IoT Security: Risks, Challenges, and Best Practices in Securing the IoT*. Digi.com; Digi International. https://www.digi.com/blog/post/iot-security?utm_source
- Rajendran, N. (2025). *Edge Computing in Finance: Real-Time Processing for Banks*. OTAVA. <https://www.otava.com/blog/edge-computing-in-finance-real-time-data-processing-for-regional-banks/>
- Terry, R. (2025). *What is Zero Trust Security? Principles of the Zero Trust Model | CrowdStrike*. Crowdstrike.com. <https://www.crowdstrike.com/en-us/cybersecurity-101/zero-trust-security/>
- Tomasz Krakowczyk. (2024). *The Role of AI and Cybersecurity in the Financial Sector*. Software Mind. <https://softwaremind.com/blog/the-role-of-ai-and-cybersecurity-in-the-financial-sector/>
- Wadhvani, K. (2024). *How Blockchain is Revolutionizing the Fintech industry. Blockchain Technology, Mobility, AI and IoT Development Company USA, Canada*. <https://doi.org/1055566/fRS0CMTS1-MCEM6OxIsC>

About the Author

Oreoluwa Joda is a Cyber Risk/Information Security manager. With years of extensive experience in the IT industry, Ore possesses an arsenal of dynamic competencies and skills spanning from her work in the Fintech industry, Education and Health Tech industry, and International Policy.

She is an academic and a three-time graduate, with a Juris Master's degree in Law, concentrating in Cybersecurity, National Security and Information Privacy, completed at George Mason University's- Antonin Scalia Law School.



Ore is also currently completing her PhD in Cybersecurity Management at Nova Southeastern University, Florida. She is a Certified AWS Cloud Practitioner, Aviatrix Multi-Cloud Network Associate, a Professional Scrum Master (PSM I), Lead Auditor Certified (ISO 27001), and has various certificates in Cybersecurity domains, and Global Financial Compliance to mention a few.

In her current position as a Cyber Risk Manager with a Global Fintech and Fortune 500 company, Ore wears a plethora of hats. She occupies a key strategic role, leading initiatives that directly influence the organization's Cybersecurity posture at scale. Her work encompasses critical domains, including Enterprise Vulnerability Management, Security Controls Assessment, Cloud and Network Security, Identity and Privileged Access Management (IAM/PAM), and Cyber Governance and Risk Compliance- integrating technical risk assessment with evolving regulatory requirements to ensure strategic alignment across business units and jurisdictions.

In addition to her Fintech and Cybersecurity experience, she also has a distinguished background in Health Technology consulting, where she was involved in the modernization of Healthcare Information Systems across the U.SA.

As a Health Tech consultant, she has traveled to over 20 U.S. states, where she trained many physicians, nurses, and healthcare administrators on the integration and effective use of Electronic Medical Records (EMR) systems- including EPIC and Cerner- within hospital and clinical operations.

Ore is dedicated to strengthening Cybersecurity resilience, regulatory compliance, and digital transformation across critical industries. Her unique interdisciplinary expertise- spanning technology, law, and policy- positions her as a strategic leader in the evolving landscape of Cyber Risk Management and Information Security. Through her work, she continues to drive impactful change at the intersection of enterprise security operations and national digital infrastructure.

Ore can be reached via email at: oj125@mynsu.nova.edu or on IG [@Ore_Global](#)



Cybersecurity Compliance – Changing the Paradigm with AI

How to leverage AI to accelerate cybersecurity compliance

By David Ramirez, CISO, Broadridge

A message to any and all financial service companies: Don't wait. Don't hesitate. Don't procrastinate. Now is the time to re-make your cybersecurity compliance leveraging artificial intelligence.

Using AI solutions broadly available from technology providers today, financial service firms can create new, modernized cybersecurity compliance systems that save time, reduce costs, lower risks and make organizations agile enough to easily adapt to future regulatory changes.

The problems with manual compliance tracking

Most financial service firms are subject to the jurisdictions of multiple regulators in their own country and also operate in other countries where regulatory requirements can differ significantly from those of their home markets. In addition, many firms must account for industry standards such as the Payment Card Industry or SWIFT's Customer Security Controls Framework (CSCF).

Until now, compliance teams have spent countless hours creating tracking mechanisms not just for each jurisdiction, but for each individual regulatory requirement within those jurisdictions. The end result is a

sprawl of documentation spelling out how the firm's controls meet each specific requirement, with evidence of controls and analysis repeated again and again to account for slight deviations in expectations from jurisdiction to jurisdiction and rule to rule.

This manual process consumes thousands of hours of time, imposes on-going demands on the control owners, cybersecurity teams and compliance specialists charged with monitoring and keeping up with regulatory changes, and creates real compliance risks for firms who make mistakes in the static documentation or fail to keep up with regulatory revisions.

The LLM solution

The emergence of Large Language Models (LLMs) is allowing financial service firms to turn that model on its head, changing what was a reactive process into a proactive risk-based program.

By applying LLMs, cybersecurity teams can create a central repository of controls and use the artificial intelligence solution to analyze and document how specific regulatory expectations are being met. The LLM can be prompted to identify the general risk being remediated by each control and then translate how that remediation applies to meet the varying expectations set by individual regulators and industry standards across jurisdictions.

Consider the example of passwords. A financial services firm operating in the United States, Europe and Asia might be subject to *at least* three regulations on password lengths:

- The EU's DORA (Digital Operational Resiliency Act) doesn't prescribe specific password lengths. Instead, it mandates that financial entities implement robust ICT risk-management frameworks, including secure authentication mechanisms, which are expected to align with industry best practices such as NIST and ENISA.
- The New York Department of Financial Services (NYDFS) Cybersecurity Regulation requires firms to implement risk-based policies for access controls, including password complexity. While the rule does not mandate a specific length, it emphasizes strong authentication and periodic review of access privileges.
- The Monetary Authority of Singapore (MAS) Technology Risk Management Guidelines recommend a minimum password length of 12 characters for privileged accounts and eight characters for standard users. The guidelines also emphasize the use of multi-factor authentication (MFA) and password expiration policies.

Rather than manually spelling out how the firm's controls meet the different expectations of each of these three rules, an AI solution uses a one-to-many approach. First, the LLM is prompted to identify the firmwide control that addresses the topic of password length. That control is then translated to the specific expectations of each requirement. The result is a report on how each and every requirement is being met by the firm's existing controls.

Let's look at another example: privileged access expectations. Here again, an international financial services company must comply with rules from the EU, the NYDFS and the MAS, among others.

Each of these regulators has its own tweaks and expectations. For example, the NYDFS explicitly requires firms to limit the number of privileged accounts, review access at least annually, and disable unnecessary accounts. It also mandates the use of MFA for privileged access and remote access. In contrast, the MAS guidelines are prescriptive: privileged access must be granted on a need-to-use basis, reviewed quarterly and logged. The EU comes down somewhere in between, requiring financial entities to implement role-based access controls and segregate duties to prevent conflicts of interest. Privileged access must be tightly controlled, monitored and periodically reviewed. DORA also mandates logging and auditing of privileged activities.

As in the password example, the LLM can be prompted to analyze the existing control and provide documentation on how that control meets the privileged access expectations set by the different regulations.

Next-gen compliance

These examples give some idea of the redundancies required in traditional manual compliance tracking. All three regulators are setting more or less the same basic requirements. However, each regulator codifies those requirements with unique expectations. These expectations vary in some ways, but overlap in many others.

An LLM-driven compliance solution will almost automatically apply the single control to each expectation. This analysis will also flag any current compliance gaps that need to be addressed, and identify gaps that will need to be addressed if there are changes to either the regulations or controls. Critically, the LLM can maintain the mapping and provide up-to-date information on mitigating controls and dependencies on a continuous basis.

These solutions can make an impact well beyond the realm of government regulation. AI solutions can also be applied to industry accreditations and certifications and other types of independent audits. For example, the Payment Card Industry, SWIFT, and ISO/IEC 27001 share a wide number of control requirements. An LLM can work from a central control repository to identify how individual requirements are met and determine if there are gaps that need to be addressed. That process will save time in both audit preparation, and during the audit itself.

By dramatically reducing the number of hours the cybersecurity team spends on documentations and audits, LLM-driven compliance approaches are reducing overhead costs. By creating a dynamic system that easily updates based on changes to internal controls or external regulatory expectations, these solutions reduce compliance risks. By bringing efficiency to the entire compliance process, artificial intelligence is freeing up cybersecurity teams to do what they are meant to do: protecting the organization from cyberthreats.

About the Author

David Ramirez is Chief Information Security Officer at Broadridge Financial Solutions, Inc. He assumed his current role in July 2024. Most recently, he was the CISO at DRW Holdings, LLC. Previously, he served as a Divisional CISO at Capital One Financial Corporation and as CISO at Brown Brothers Harriman & Co. Earlier in his career, he held positions at Bridgewater Associates, Barclays PLC, Alcatel-Lucent, DVS International, Intelsys and Cyberia Group. Ramirez earned an MS in Information Technology Security from the University of Westminster and a Diploma in Risk Management from the Institute of Risk Management.





Cybersecurity In Critical Infrastructure: Protecting Power Grids and Smart Grids

By Kehinde Ayano, Assistant Professor of Computer and information Science, Indiana Wesleyan University Marion Indiana USA

Infrastructure like water system, supply system, telecommunication networks, and power plants are critical assets for any country in that the destruction and incapacity of such systems poses an adverse effect on security, economy, health, and overall welfare and existence of any country. The integration of digital and cyber warfare into traditional warfare has necessitated the need to adequately secure those critical infrastructures as they become top target by state actors in case of conflict and war.

Most systems in modern society are electricity driven which makes power and smart grids very crucial as they underpin nearly all other critical infrastructure. A successful attack on this infrastructure will have a cascading effect on all other critical infrastructures. This article discusses the evolution of power grids, threat landscape and vulnerabilities in power and smart grids. It also examines real world case studies of cyber- attacks on power and smart grids analyzing the incidents and concludes with security strategies and best practices for protecting power and smart grids.

Evolution of Power and Smart Grids

Traditional power system also known as power grids are a one-way system for distribution of electricity from producers to consumers and are vital for functioning of businesses, society, and government at large. They are manually controlled with limited capacity for integration with renewable energy. Advance in technology and digital evolution led to the development of modern versions of the traditional power system that makes use of digital technologies for monitoring, management, synchronization, and transportation of energy from multiple sources to meet the varying demands of the consumers. These smart grids, unlike the power grids, are two-way communication systems with automated control and real time monitoring and allows for easy integration of renewable energy which improves the reliability and efficiency of electrical power systems.

Components of Smart Grids Communication Network

Some of the major components of the smart grid communication network include the following which allows for seamless two-way communication between utilities and consumers include the following.

Control Center: This is the central hub for monitoring and managing the entire grid. It accepts data from all other components and sends control signals for grid operation management.

Substation: Transforms high voltage from the transmission network to lower levels suitable for distribution. Smart grids substations are equipped with sensors and devices that can send data on power quality, load condition and status of equipment to the control center.

Smart Meter: Smart meter measures and communicates consumption with both consumer and the utility in real time.

Advanced Metering Infrastructure: It facilitates communications between smart meters and utilities, and send smart meter data to the control center and other grid components

The components listed above and many more make smart grids a fully digitalized communication network improve reliability and efficiency of electrical power system. However, the integration of digital technology in smart grids also introduces new vulnerabilities and cybersecurity threats that must be addressed for robust operation. Ensuring that power and smart grids are secured is critical to the existence of business, organization, and government as the resultant of these attacks could be catastrophic and life threatening.

Threat landscape in Power and Smart Grids

Malware: These are malicious software designed to disrupt damage and gain access to the system. This includes trojans, virus, ransomware, and many others. Malware exploits known and zero-day vulnerabilities in software, hardware and network protocols used in power systems and can disable or disrupt Supervisory control and Data Acquisition systems SCADA, DCS and other operational technologies.

Phishing: This is a form of attack whereby an attacker disguises and attempt to acquire sensitive information such as usernames and passwords by posing to be a legitimate entity.

Network Intrusion: Network communication systems of power and smart grids can be intruded through weak security configurations like default password, unsecured remote access, or unpatched systems and other vulnerabilities to gain control into the system.

Distributed Denial of Service (DDOS): This is an attempt to disrupt the availability of services provided by smart grids and make them unavailable by overwhelming the system with traffic from multiple sources. The DDOS are usually launched from malware infected hosts and could be volume-based attacks like UDP and ICMP floods, protocol attacks like SYN flood and Smurf DDOS or Application layer attack GET/POST floods.

Advanced Persistent Threats (APT): This is a prolonged and targeted cyber-attack whereby state actors or highly skilled cyber criminals gain access to a network and remain s undetected for an extended period.

Vulnerabilities in Power and Smart Grids

The attack surface has significantly expanded in smart grids due to complex network of devices which includes sensors, smart meters, smart switches, communication networks and control systems with each of these components being a target for cyber-attacks. Increased connectivity and data exchange within the control center and other components of smart grids make it more vulnerable to attack. Therefore, to maintain the resilience and security of smart grids, understanding and addressing the vulnerabilities inherent in smart grids systems is critical.

These vulnerabilities include the following:

Legacy Systems: The continuous use of Legacy systems which are outdated technologies due to certain constraints within an organization, poses significant risk to the security of such systems. This is because such systems may no longer be patched for updates and may also have limited monitoring capability.

Interconnected Networks: The vast interconnection of devices and increased connectivity of communication systems of smart grids if not properly secured, make them highly vulnerable to attack.

Remote Access: The management and monitoring of grids system are usually done through remote access. Vulnerabilities in remote access connection may be exploited by attackers to gain access into the system.

Supply Chain Risk: Smart grids heavily rely on complex supply chain of hardware and software components which are majorly contracted out to manufacturers and suppliers. The security practices of such 3rd party vendors, if not robust, may pose significant risk when integrated into the power and smart grids. Attackers can also target the software development lifecycle by compromising legitimate software and software updates which in turn makes the system in which they are deployed vulnerable to attack. An example of such supply chain vulnerabilities is the SolarWinds attack (2020) where malware is injected into routine software update.

Human Factor: Human factor is one of the most common vulnerabilities in cybersecurity framework. Error and negligence or malicious intent by staff despite the solid technological defenses have led to system compromise. This compromise comes because of inadequate training and awareness, poor password practices and insider threats.

Real World Examples of Cyber-Attacks on Power and Smart Grids

Due to the digital evolution of electrical power systems, power and smart grids are increasingly becoming ground zero for cyberwarfare. Over the past two decades, several attacks have been launched against smart grids resulting in outages and financial loss resulting from payment of huge ransom. Example of such is the attack on Ukraine Power Grid in 2015 in which BlackEnergy malware was used to compromise three Ukrainian distribution system using spear-phishing email. The attacker gained access to the Supervisory Control and Data Acquisition (SCADA) systems and compromised the circuit breaker remotely and disabled the UPS and Backup. Also, in 2016, Ukrainian transmission station was targeted by a custom-built malware named Industroyer which compromised the Industrial Control System and disrupt power distribution for about an hour.

Correction from the previously published version of this article (August 2024): An earlier version of this article stated that Florida municipal power agencies were targeted in June 2021 using phishing and remote vulnerabilities. This statement could not be substantiated and has been removed to ensure accuracy. Nonetheless, these cases underscore the importance of security strategies and best practices in power and smart grids management.

Security Strategies and Best Practices for Managing Power and Smart Grids

Cyberattacks on power grids and smart grids have become more frequent and sophisticated in recent years and can have devastating consequences which include blackouts, economic losses, disruptions to vital infrastructure, and theft of sensitive data. Therefore, there is a need to put in place sound security strategies and best practices to safeguard this critical infrastructure from attack. Some security strategies and best practices for power and smart grids are discussed below.

Risk assessment and management: Risk assessment and management plays a vital role in the security of power and smart grids as they help to detect and mitigate vulnerabilities and help in incidence response. Implementing Risk assessment and management using the NIST **Interagency Report (IR) 7628 Revision 1** which provides a comprehensive framework for securing smart grid systems will go a long way in securing this critical infrastructure.

Defense-in-Depth: Implementing a layered security approach using various security controls and protocols (firewalls, encryption, IDS, IPS, SIEM, access controls) will enhance the security posture of smart grid systems.

Vulnerability Assessment and Penetration Testing: Detecting inherent weakness in smart grid systems before an attacker does through comprehensive vulnerability assessment and simulation of real attack to discover vulnerabilities that are hidden and remain undiscovered by automated scanning will allow those security lapses in the system to be tightened before they are exploited on by attackers.

Patch Management: Apart from ensuring system reliability, effective patch management also reduces attack surface. It is more cost-effective to proactively address vulnerabilities in smart grid through effective patch management than to reactively mitigate the resultant effect of security breaches.

Network Segmentation: Segmentation of communications network system of a smart grid system inhibits lateral movement preventing attacker from gaining access to the entire system in case of breach thereby minimizing the impact of the attack. It also helps remediation as focus can be only on the compromised segment.

Data Backup and Recovery (BCP and DRP) Plan: Having a Business Continuity Plan and Disaster Recovery Plan in place will help to facilitate recovery from cyber-attacks, reducing time and mitigate the impact on services.

Employee Training and Awareness Programs: The importance of employee training and awareness could not be overemphasized as research has shown that humans are the missing link in the cybersecurity chain as they are highly susceptible to social engineering, phishing, insider threats and prone to commit errors. Training and awareness will help employees to have good cyber hygiene and cultivate strong cybersecurity structure.

Conclusion

In conclusion, power and smart grids security requires a multidimensional approach that combines implementation of security controls which are administrative, physical, and technological, and proactive risk assessment and management, and continuous training and retraining of human elements. Making cybersecurity a top priority and fostering cybersecurity culture will safeguard this critical infrastructure from attacks.

About the Author

Kehinde Ayano Ph. D. is an assistant professor of Computer and Information Science at Indiana Wesleyan University Indiana. He is also a Certified Information System Security Specialist. Kenny can be reached on Kenny.ayano@indwes.edu.





Data Breaches Are the New Normal – Complacency Is the Real Crisis

By Rob Vann, Chief Solutions Officer at Cyberfort

When a company like Qantas an airline synonymous with safety suffers a high-profile data breach, the message is loud and clear: no brand is untouchable, and no data is sacred. But here's the real problem: we're still treating breaches as anomalies. They're not. Breaches are now a guarantee, and the only variable left is how well or how catastrophically you respond.

The Qantas breach wasn't just a failure of security; it was a failure of imagination, of preparation and resilience. If businesses don't wake up now, they won't just lose customer trust they'll lose relevance. This is your blueprint for what to do when, not if, your defences fail and how to ensure your organisation doesn't become the next cautionary headline.

Step one: Panic smart – not fast

When the breach hits, most companies do the same thing: go silent, scramble internally, and throw together a press statement that says, "We take your privacy seriously."

Stop. That's PR autopilot and attackers are counting on it.

What you need is speed with clarity. Assemble your breach response team legal, security, comms, compliance and ask the hard questions:

- What exactly was accessed?
- How long has it been going on?
- Is the attacker still inside?

The longer you pretend it's "under investigation," the more trust you lose. Transparency isn't just a legal risk it's a strategic advantage.

Consumers don't wait to be told

If you're a Qantas customer (or one of the millions watching nervously), don't sit around for confirmation. Assume compromise until proven otherwise. Cybercriminals won't wait for your email to arrive they'll be monetising your data by tomorrow.

Verify the breach – don't fall for the follow-up scam

Ironically, the breach itself often triggers a second wave of fraud. Phishing emails pretending to be from Qantas will flood inboxes, asking you to "verify your account" or "reset your details." Never click on email links after a breach. Go directly to the company's website or app. Trust your paranoia it might save your identity.

Check if you've been exposed – and act accordingly

Not all data breaches are created equal. A leaked email is annoying. A leaked passport number? That's catastrophic.

- Use monitoring tools like HaveIBeenPwned or sign up for dark web scanning through your bank or a cybersecurity provider.
- For loyalty and travel accounts, scrutinise redemption histories and account logins. Flag anything out of pattern.
- If ID documents were leaked, report them immediately and request replacements or fraud alerts with the relevant authorities.

The attackers won't give you time to think. Don't give them time to act.

Password resetting isn't optional. It's urgent.

Still using the same password you created in 2012? Then you're part of the problem.

Qantas frequent flyer accounts are a prime target because people reuse those passwords everywhere - banking, email, e-commerce. One breach becomes many.

Your new password rulebook:

- Unique for every site.
- Long (at least 12 characters).
- Random (not "Qantas123!" or your child's name).
- Managed with a password manager. You don't have to remember 100 passwords – you just need to remember one good one.

Weak passwords don't get guessed, they get cracked by bots running billions of combinations in seconds. If you're still relying on "clever" variations, you're already compromised.

Two factor authentication isn't a luxury. It's a minimum requirement

Two-Factor Authentication (2FA) is one of the simplest, most effective ways to stop account takeovers. So why aren't more people using it?

Excuses like "it's annoying" or "I don't want to install another app" don't hold up when your identity is at risk.

Here's what to do:

- Enable 2FA on every account that offers it—especially loyalty programmes, email, and banking.
- Use an authenticator app (like Microsoft or Google Authenticator) -NOT SMS, which is easier to hijack.
- Never share or screenshot your authentication codes. They're like handing out keys to your digital kingdom.
- Shop and travel smarter: Assume you're being watched
- Cybercriminals love predictable behaviours. Travel is full of them.
- People use unsecured Wi-Fi in airports and hotels.
- They receive dozens of emails from travel brands.
- They're often distracted, tired, or rushed -perfect conditions for phishing.
- Consumer Tips:
 - Don't shop or log in to sensitive accounts over public Wi-Fi unless you're using a VPN.
 - Never use the same email/password combo across shopping and travel sites.
 - Use disposable or virtual cards when booking trips or buying online.
 - Set up bank alerts for any purchase or login activity.
 - Treat every digital interaction while travelling like it's under surveillance—because it probably is.

For businesses: prevention is dead. Resilience is everything.

Still thinking cyber “won’t happen to us”? Ask Qantas. Ask MOVEit. Ask anyone who’s had to face the cameras and say, “We’re investigating the incident.”

You don’t stop breaches with wishful thinking and legacy tools. You stop them with brutally honest assessments, relentless testing, and round-the-clock visibility. Three key steps all organisations should be taking in light of the Qantas breach:

1. Penetration testing – Simulate the breach before the real one hits

Static security reviews are useless in 2025. Attackers don’t use checklists, they use ingenuity. Your defences should be tested by people who think like them.

Use red teams to run real-world attack simulations to expose your blind spots, from credential stuffing to insider threats. If your internal team always passes the test, it’s not a test. It’s theatre.

2. Managed detection & response (MDR) – Eyes on everything, all the time

Breaches don’t announce themselves. Without MDR, you might not know you’ve been hit until your data is on the dark web. Market leading MDR platforms use AI to detect anomalies in real time, and expert analysts investigate alerts before they become incidents. Speed matters. Context matters more. If you’re relying on tools alone, you’re not covered, you’re exposed.

3. Secure cloud backups – Because ransomware doesn’t negotiate

When all else fails, your backup is your survival plan. But if it’s stored on the same network, with the same credentials, and hasn’t been tested in six months, you might as well not have one.

A proper backup strategy includes:

- Isolated, encrypted cloud storage
- Automated versioning
- Disaster recovery plans that are rehearsed, not theoretical

If your board doesn’t know your RTO (Recovery Time Objective), ask why they still have a seat at the table.

Final word: The real breach is the illusion of control

Let’s stop pretending we can “prevent” all cyber-attacks. That ship has sailed. What separates survivors from casualties is preparedness, transparency, and relentless resilience. Qantas didn’t choose to be breached, but they did have a choice in how ready they were when it happened.

For consumers - assume you’ve been compromised and act accordingly. For businesses - build breach response into your DNA.

This isn't about fear. It's about facing reality. Cyberattacks are business attacks, and the cost of not evolving is far greater than the cost of change.

Because in today's world, data protection isn't just a duty, it's your credibility.

About the Author

Rob Vann is the Chief Solutions Officer at Cyberfort, bringing in over 35 years of experience in cloud security and managed services. He has successfully led cybersecurity initiatives across various industries, including large corporations, telecommunications, and government sectors. Vann has played a key role in developing managed security services, scaling businesses to revenues exceeding £50 million per year.

Rob can be reached online via [LinkedIn](#) and at our company website <https://cyberfortgroup.com/>





Deepfakes and Disinformation: Protecting Truth in the Age of AI

By Omkar Bhalekar, Sr Network Engineer, Tesla

In today's digital era, borders between truth and fiction are fading. As artificial intelligence technologies emerged, the stage was set for a new and malevolent danger to arise, deepfakes. Deepfakes are fake media produced with the help of artificial intelligence that can create extremely realistic images, audio, and video, which have spawned an ever-growing wave of disinformation dismantling our very definition of what truth is. This blog discusses the threat of deepfakes, their use in disseminating fake news, and the actions we need to take to safeguard truth in this age of AI.

What Are Deepfakes?

They are created by advanced machine learning algorithms called Generative Adversarial Networks (GANs). Networks are founded on a competition between two artificial intelligence actors: one creates

fake material and the other attempts to recognize if material is actual or manipulated. Deepfakes are extremely realistic by the adversarial technique, and even experts have difficulty distinguishing manipulated media from actual content.

Whether a politician accused of saying things he never said, or a voice-over impersonating someone to confirm dubious payments, deepfakes are the new threat of information manipulation.



[Tom Cruise Impersonator](#)

The New Disinformation Threat

Disinformation is nothing new, propaganda and disinformation have been around for thousands of years. But deepfakes represent a new devilish twist in the fact that they make spurious information both visually and audibly plausible, multiplying their impact exponentially.

Eroding Public Confidence

While doctored clips go in all directions, they create confusion, mistrust, and suspicion even on true facts. Such a corrosion of trust undermines the foundations of democratic governments and institutions and facilitates evil forces to shape the opinion of the public and societal fragmentation.

Manipulating Politics and Elections

Political deepfakes can determine the outcome of an election by spreading false information or discrediting a candidate. Consider a doctored video of a politician making inflammatory speeches a week before an election, such an action would disenfranchise undecided voters or lead to rioting.

Inciting Social Discord

The deepfakes could be used to generate social instability in the form of rumors against ethnic minorities, religious groups or social movements. This polarizes, result in protest, or riots.

Facilitating Identity Theft and Fraud

Deepfake video and audio are increasingly being utilized by attackers in newer and more advanced methods to impersonate CEOs or government officials and manipulate employees or citizens into making fund transfers or sharing sensitive information. The attacks can be extremely targeted and bypass traditional security practices like passwords or security questions.

Examples of Deepfake Misuse in Real Life

1. The Nigerian Email Scam Using Deepfake Audio

In 2019, the CEO of a UK energy firm was [deep faked](#) via voice to direct an employee to wire €220,000 to a Hungarian supplier. The accent employed was the CEO's German one and sounded very natural and led to a very expensive scam.

2. Ukrainian President Zelenskyy Political Deepfake

[Doctored videos](#) cropped up in 2022 amid the Russian invasion of Ukraine, falsely with Ukrainian President Volodymyr Zelenskyy speaking words of surrender. Doctored videos were intended to demoralize the Ukrainians and mislead global audiences.

3. Deepfake Porn and Celebrity Exploitation

There are a number of hundred celebrities and [public figures](#) who have been attacked by [deepfake porn](#), with celebrity faces superimposed on adult content via AI without authorization. It is not only an invasion of privacy but also causes harassment and defamation.

4. Fake Biden Speech Video (2023)

One of the earliest high-profile deepfakes to go viral was of US President Joe Biden in 2023, slurring his speech. Although it had been edited earlier and created [misinformation](#) on the internet, people were confused and started commenting politically.

5. AI-Generated Deepfake CEO Scam (2024)

Recent instances have shown how extensively the technology of deepfakes has been used by hackers to [impersonate CEOs](#) in video calls and blackmail workers into approving huge transactions within seconds. Such hacks are a chilling level of [sophistication](#).

How can we Detect and Combat Deepfakes?

Deepfakes war is multi-faceted and will have to be addressed with technology, education, and policy.

- **AI-Based Detection Tools**

Inventive technology companies and researchers have developed AI programs that monitor video and audio for characteristic signs of lying, i.e. inconsistent blinking, strange facial response, or audio glitch. Technology like Microsoft's Video Authenticator can attach a confidence level to content media, allowing individuals to verify as authentic before uploading or acting on it.

- **Digital Watermarking and Blockchain**

New technologies are incorporating invisible digital watermarks in material or keeping metadata for blockchain transactions to enable the content to be validated. This suggests that content creators and media companies will be able to determine if their material is original or not, thus making it simple to detect imitation.

- **Media Literacy and Public Awareness**

The best defense is an educated public. Media literacy educates people to recognize when they are being deceived and to ask who the source is, and to go check it out before you share it. Governments, teachers and platforms all have a role to play in promoting critical thinking on the web.

- **Policy and Legal Frameworks**

Various nations are drafting regulations to criminalize the malicious use of deepfakes but also safeguard free speech and creativity. California has now defined non-consensual deepfake pornography as a crime while the EU is mulling over rules to curb the proliferation of AI-generated content.

- **Industry Cooperation and Standards**

Technology companies, media organizations, and governments all must collaborate to establish content verification systems and rapid response to deepfakes. Shared databases and transparency reports of all previously known fakes will be able to prevent the spread and impact.

The Ethical and Societal Challenge

Deepfakes are not inherently evil; they also represent creative possibilities in film, game, and access, i.e. natural-sounding dubbing or bringing dead people back. The ethics problem is balance between innovation and responsibility. With deepfake technologies that are becoming more advanced, individuals and communities will have to deal with new issues of consent, privacy, and what truth means.

Staying One Step Ahead

Deepfake tech is still accelerating at breakneck pace, making detection and defense an endless cat-and-mouse game. But through innovative thinking, regulation, education, and awareness, we can maintain the risks in proportion and utilize AI's promise for good.

Defense of truth in the age of AI isn't a tech problem but it's everybody's problem that requires cooperation between disciplines and borders.


Conclusion

Deepfakes are an ancient adversary of the AI age, a potentially disastrous capability to be used for fantasizing or letting go. Managing reality requires attention, technological progress, legislating, and an educated population. By disseminating awareness of the danger and enforcing coercive countermeasures, we can ensure truth prevails over more man-made existence.


In the algorithmic age, reality is fluid and when AI distorts the mirror, only truth shows what is real.

About the Author

Omkar Bhalekar is a Sr Network Engineer at Tesla Motors. He specializes in advanced networking in the field of smart and sustainable manufacturing technologies of electric Vehicles, battery technology, and Robotics following Industry 5.0 standards which focusses on environmental sustainability. With 6+ years' experience in this field and cybersecurity enthusiast specializing in Data center architecture, Manufacturing infrastructure, and Sustainable solutions. With extensive experience in designing and securing resilient industrial networks, building smart factories and AI data centers with scalable networks, Omkar avidly writes to simplify complex technical topics for engineers, researchers, and industry leaders.



Omkar can be reached online at LinkedIn: <https://www.linkedin.com/in/omkar-bhalekar> or at his personal website <https://www.omkarbhalekar.com/>



Securing the Mission: Why External Cyber Defense is Essential for Government and Education

By Amit Weigman, Office of the CTO, Check Point Software Technologies

In the past few years, ransomware crews and nation-state hackers have turned America's state agencies, city halls, and school systems into prime hunting grounds. Budget-strapped information-security teams that once weathered the occasional intrusion now face a relentless barrage of extortion, data theft, and disruptive campaigns. The shift is no accident; wars, realignments in global politics, and the industrialization of cyber crime have **lowered the technical and financial barriers** to launching sophisticated attacks. Criminal syndicates freely purchase tooling once reserved for intelligence services, while state-sponsored groups borrow from the playbooks of profit-driven hackers. For the public sector, especially the State, Local, and Education (SLED) vertical, the threat surface has never been broader or more complex.

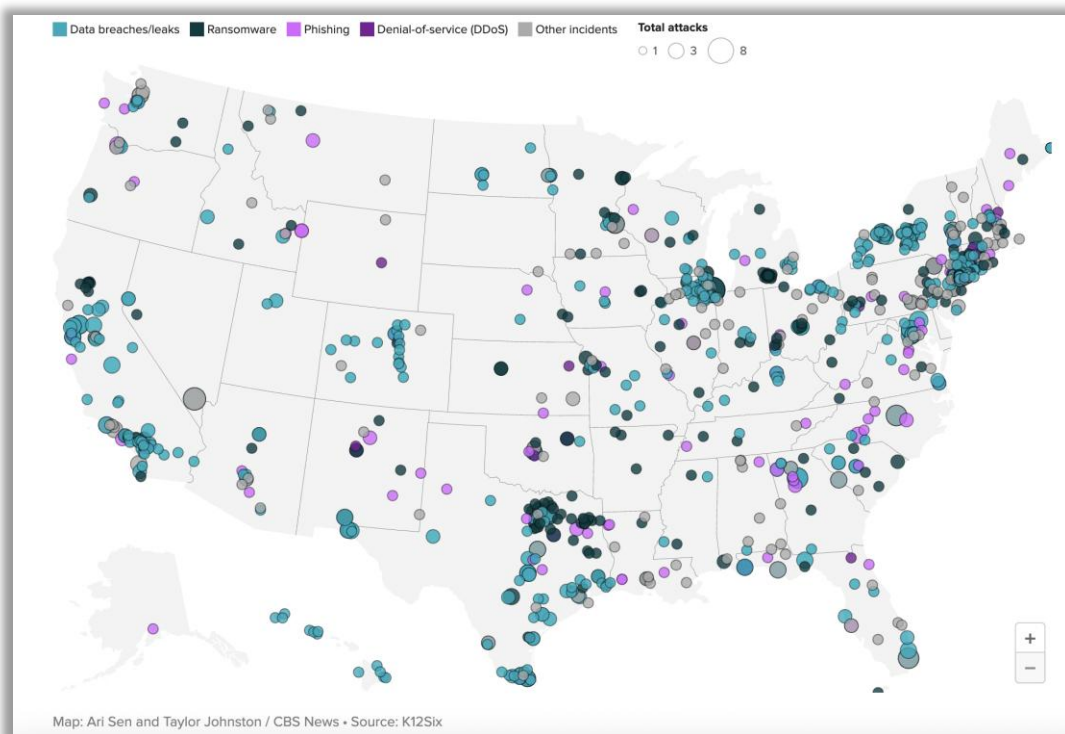


Figure 1: Cyber Threat Landscape Across U.S. Public Sector Entities

Traditional security cannot keep pace with this reality. Detecting malware already running inside the network or scrambling to patch a zero-day after it has been weaponized amounts to fighting yesterday's war. What SLED organizations need is external cyber defense: **continuous visibility into the attacker's ecosystem** so threats can be seen, understood, and disrupted before they ever reach official infrastructure. Below are six pillars of an external-first approach that form a blueprint for safeguarding public services and the citizens who rely on them.

1. Deep and Dark Web Monitoring

The underground economy has matured into a vast marketplace where credentials, network access, and confidential documents are traded by the terabyte. School districts and state health departments can't afford to overlook those markets. A single infostealer infection on a student laptop can harvest credentials that, weeks later, appear for sale to ransomware affiliates. Timely identification of those stolen logins turns a looming crisis into a simple password-reset exercise. Continuous dark-web monitoring serves as an early-warning system; it spots the early steps of an attack while defenders still have room to act.

2. Phishing and Brand-Impersonation Detection

Recent campaigns have impersonated departments of motor vehicles, for example, sending citizens fraudulent text messages about unpaid tolls and directing them to cloned payment portals. Such schemes harvest credentials, steal funds, and erode confidence in government services. External monitoring that

hunts for suspicious domains, look-alike social-media pages, and rogue email infrastructure can uncover these operations within hours of setup, often well before the first phishing message lands in an inbox. Rapid takedown of malicious assets disrupts the attack chain and protects both the agency and the public.

3. Fraud and Financial-Crime Intelligence

While ransomware headlines dominate news cycles, quieter forms of cyber-enabled fraud siphon away staggering sums. Underground forums advertise forged passports, birth certificates, and driver's licenses that can be used to create synthetic identities or defraud benefits programs, among a wide variety of other services. By tracing these illicit services and mapping their connections to specific schemes, analysts can alert victim agencies, cancel fraudulent transactions, and strengthen identity-verification processes before losses escalate.

4. Threat-Actor Profiling

Knowing an adversary's preferred tools and tactics turns abstract risk into actionable defense. Overlaying fresh incident data such as new malware samples, server infrastructure, or social-media chatter onto those profiles helps defenders predict the next move and pre-position controls. For instance, if a group that typically exploits remote-desktop services begins experimenting with phishing kits, an agency can accelerate email-security projects or staff-training initiatives before the tactic matures.

5. Supply-Chain Risk Management

A government network is only as secure as its vendors. When a major cloud provider suffered a breach earlier this year, troves of credentials stored in customer environments, many belonging to municipalities and public universities, were exposed long before on-premises sensors fired an alert. External intelligence that monitors third-party security posture and flags chatter about newly exploited vulnerabilities narrows the gap between compromise and containment. Agencies can quickly inventory where they use an affected service, rotate keys, or isolate integrations rather than discovering exposure weeks later through an incident-response report.

6. Executive-Protection and Geopolitical Alerts

School-board presidents, police chiefs, and state CIOs increasingly find themselves personally targeted by disinformation, deepfakes, or doxxing campaigns. Fake social-media profiles that mimic public figures can solicit donations, spread malware, or inflame political tensions. Simultaneously, state-sponsored actors time cyber operations to coincide with local events such as election primaries or contract negotiations to maximize disruption. Alerts that fuse geopolitical context with technical indicators allow agencies to prepare for region-specific threats, coordinate with partners, and communicate clearly with constituents when false narratives surface.

Building a Proactive Posture

Implementing these pillars starts with culture. Security teams must expand their field of vision from “inside-out” to “outside-in,” treating open-source intelligence, criminal marketplaces, and third-party ecosystems as integral components of the attack surface. Practically, that means deploying external sensors and crawlers that collect data from forums, paste sites, phishing kits, and vulnerability disclosures in real time; automating enrichment and correlation so stolen credentials, malicious domains, and threat-actor chatter surface in the same console analysts use for internal alerts; integrating intelligence into workflows, from instant ticket creation when employee passwords appear for sale to automated blocking of typosquatted domains at the web proxy; and measuring success by dwell time outside the perimeter: the sooner stolen data is discovered or an impersonating site is dismantled, the less time attackers have to weaponize it.

Collaboration is equally important. Local agencies benefit from sharing Indicators of Compromise and playbooks with neighboring counties; universities gain insight by feeding anonymized data into sector-wide Information Sharing and Analysis Centers. When an attacker reuses infrastructure across multiple victims, cross-organizational intelligence lets defenders cut off entire campaigns rather than fight piecemeal skirmishes.

The Road Ahead

Despite the intense and changing threat landscape, the path to resilience for SLED organizations is becoming clearer; external cyber defense shifts the advantage back to defenders. Agencies that invest **now** in programs that marry dark-web telemetry, phishing takedown, fraud analytics, and third-party risk monitoring will not simply react faster; **they will force adversaries to work harder, spend more, and accept a shrinking return on each attempted intrusion.**

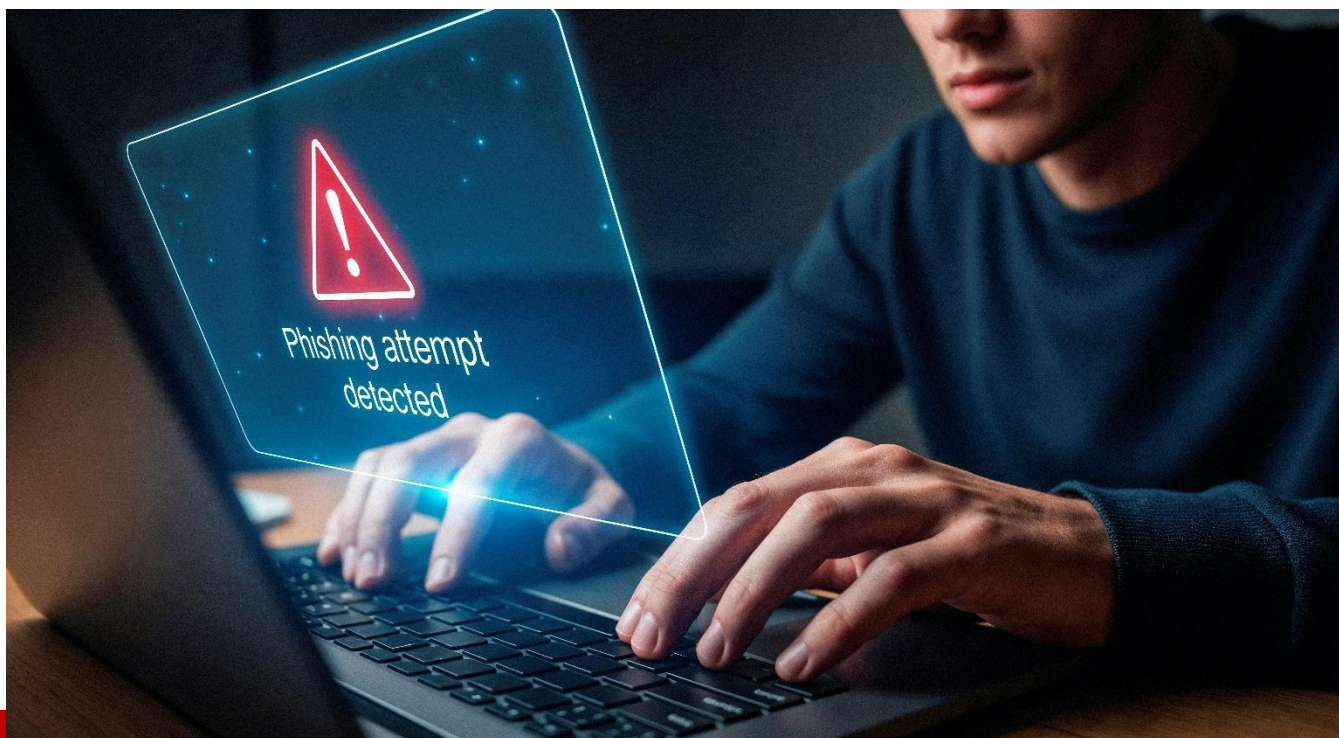
About the Author

Amit Weigman is an experienced Project Specialist and Cyber Security Analyst with a demonstrated history of working in the military industry. Skilled in Intelligence Analysis and Cyber Threat and Security Analysis, as well as Molecular Biology.

LinkedIn profile: <https://www.linkedin.com/in/jaredvichengrad/>

Company page: <http://www.checkpoint.com/>





Erase, Don't Just Delete: How Secure File Wiping Can Stop Insider Threats

By Neha Sawhney, Marketing Specialist, Stellar Information Technology Pvt. Ltd.

Imagine you have just cleaned out the recycle bin and feel like you have waved a magic wand at those embarrassing files of yours and you start to think they are gone for good but that is not entirely the case. Hitting that **Delete** button on a PC is more like shoving your messy papers under the rug... they are still there, yet they are just out of sight. In truth, these deleted files normally keep lingering on the drive until the time someone or something overwrites them and a determined insider (or maybe a snooping coworker with a forensic toolset) can go ahead and uncover them with a few clicks. It's a classic bait-and-switch. When you think the data vanished, it's really just "sleeping" on the disk. For anyone serious about security, this is a nerve-wracking illusion. and a reminder that we need to **erase**, not just delete.

Understanding the Insider Threat

An insider threat originates from individuals within your organization or those you are already familiar with. This could include current or former employees, contractors, or business partners who possess legitimate access to systems and data—but misuse that access, intentionally or unintentionally, to compromise security or violate privacy. In other words, this isn't a hooded hacker breaking in from the

internet; it's someone *you already let in the front door*. These threats are surprisingly common – in fact, a 2024 [survey](#) found that **83% of organizations** experienced at least one insider attack in the past year. Why so many? Because insiders have a distinct advantage, they know the systems, the people, and the weak spots, making their malicious moves hard to spot.

In most cases, someone on the *inside* either hands over or left behind sensitive info. In addition, because traditional security often focuses on external foes, insiders can slip right past standard defenses. It's a wake-up call: to protect data, companies must guard against threats coming from within, not just outside.

The Illusion of Deletion

For many, the “Delete” key feels like a data-stomping out command. However, in reality, file deletion on computers is mostly cosmetic. When you delete a file (and even empty the Recycle Bin), the operating system merely marks that disk space as “free”. It removes the file's entry in the directory. The actual bits of your file remain on the disk until something new overwrites them. Security experts put it bluntly, “simply dragging [files] to the trash bin isn't enough” because the data stays put until overwritten. As one Windows insider forum contributor explained, deleting a file *“makes it invisible and inaccessible through normal means,”* but *“the actual data remains on the disk until it's overwritten by new data.”*

In practice, this means deleted files can be recovered. Specialized recovery tools or even forensic services can reconstruct your “gone” file from the non-overwritten sectors. Even something as simple as a leftover spreadsheet or photo can be brought back with a quick scan. The idea that deletion equals destruction is a dangerous myth.

Secure File Erasure as a Defense Strategy

Secure file erasure—also known as data shredding or sanitization—serves as a reliable solution to the false sense of security that comes with simply deleting files. Rather than merely relocating data to unallocated space, secure file erasers overwrite the original information, often multiple times, using specific patterns to ensure it cannot be recovered. Good erasure tools follow rigorous standards. For instance, some methods overwrite data **3 to 7 times** with layers of random bits (finishing with a pass of zeros) to thwart any remaining trace.

Secure erasure complements other security measures. It sits at the very end of the data lifecycle: after you've transferred or backed up what you need, you run the eraser so nothing legible remains. Sometimes this involves built-in commands (like Windows' cipher /w or diskpart clean) or bootable wipes like DBAN. For SSDs, you might use hardware-level secure erase or encrypt-then-throw-away-keys, since wear leveling complicates ordinary overwrites. After erasure, if an insider (current or former) tries to grab data from a retired disk or hidden folder, there's nothing readable there. In effect, secure wiping turns potential leaks into dead ends. **Data thieves simply won't find files to recover.**

Secure erasure also helps with **compliance**. Regulations like GDPR, HIPAA, or CCPA require organizations to securely delete personal data and sometimes to prove that deletion.

Introducing Stellar File Eraser

If you're convinced that deleting isn't enough, what's a good tool to actually *wipe* data? One solid choice is **Stellar File Eraser**, a PC utility designed to permanently erase sensitive files and traces. Notably, it's available as a Free File Eraser Software (**with a full-feature 30-day free trial**), so anyone can try it without cost.

Stellar File Eraser's core engine is built for thoroughness: it uses advanced, multi-pass algorithms (including DoD-style standards) to overwrite data beyond recovery. You can target individual files/folders, entire drives, or even wipe internet traces and temporary files. It **supports a wide range of devices**... hard drives, SSDs, USB sticks, memory cards, and even network servers.

Key features include:

- **Free & user-friendly:** It advertises itself as a *free file erasure tool* with a clean, intuitive interface. There is no steep learning curve, so even non-experts can use it.
- **Multiple erasure methods:** You can choose Quick Erase, Secure Erase, or Custom Erase modes to balance speed vs. thoroughness. For casual cleanup, one pass might suffice; for highly sensitive data, you can opt for deeper multi-pass wipes.
- **Scheduling/Automation:** The software can run tasks on a schedule (one-time, daily, weekly, monthly, or at system startup) so that wiping happens automatically. This "set it and forget it" approach is great for routine clean-ups (like weekly internet history wipes or monthly archive purges).
- **Comprehensive cleanup:** It will delete not just documents and images, but also browsing history, saved passwords, application traces, and even wipe unused disk space. In other words, it is not just a file eraser – it is a *digital privacy sweeper*.
- **Detailed reports:** After each erase operation, Stellar generates a report (PDF/XML) that logs what was wiped. This is handy for compliance or audit purposes, showing *proof* that data was destroyed.

In sum, **Stellar File Eraser** is a reliable, user-friendly solution for secure deletion on Windows. And best of all, during its trial period it's essentially free file eraser software, so teams can evaluate its power without upfront cost.

Best Practices for Implementation

To make the most of secure wiping (and to actually stop insider leaks), follow some practical tips:

- **Back up first:** Before any irreversible wipe, ensure important data is backed up elsewhere. It sounds basic, but it can't be overstated. You *will* sometimes delete the wrong file.
- **Layer your methods:** Don't rely on a single action. A solid approach is: 1) delete files normally and empty the bin, 2) overwrite free space, and 3) use a dedicated eraser like Stellar for the final pass. For extremely sensitive info, consider encrypting it first and then wiping the encryption keys.

- **Use scheduling and automation:** Don't wait for the next data breach scare to remember to wipe. Set up your eraser tool on a regular schedule. Stellar File Eraser and similar tools support automated jobs, which reduces the chance of human forgetfulness.
- **Cover all copies (and clouds):** Remember, data often has backups or cloud copies lurking. Deleting a file on one computer doesn't automatically erase it from Google Drive, Dropbox, or backup servers. So make sure your wipe policy includes removing files from cloud services and backups if needed
- **Document and audit:** Incorporate wiping into your official security policies. Keep logs of wipe actions and use the eraser's reports as evidence. This not only helps with audits, but also ensures someone is actually *responsible* for data sanitization.
- **Train and enforce:** Finally, educate your team. People tend to underestimate how easy data recovery is. Teach users why they should use secure delete tools (not just the Recycle Bin) and enforce rules for decommissioning PCs, retiring drives, or off-boarding employees. Combine wiping with good access controls and monitoring (even behavior analytics) for a layered defense.

Do this regularly, and you can significantly reduce the risk of a careless mistake or a malicious act causing irreversible damage to your privacy.

Conclusion

In a nutshell, **deletion alone is a dangerous illusion**. Data only truly vanishes when it's been overwritten or destroyed beyond recovery. For IT pros and security teams, that means treating the delete key as the *first* step, not the last. Secure file wiping is a powerful layer of defense against insider threats: it makes sure that even if someone *can* access an old drive or account, there's nothing of value to take. Tools like Stellar File Eraser (yes, even in its free form) put military-grade wiping technology into your hands. Pairing such tools with good policies, such as backups, encryption, documentation, and user training, ensures that when it's time for data to go, it really **goes**. After all, once data "disappears" into an insider's hands, no amount of deleting can undo the damage. Better to slash it into unreadable bits first and truly say **goodbye** for good.



From Shadows to Spotlight—Why Low-Priority Assets Now Define the Frontline of Cybercrime

By Pankit Desai, CEO of Sequaretek

Every fortified network guards its ramparts and main gates with disposition and purpose. Yet, the most notorious intruders throughout history seldom opted for direct assaults. Instead, their cunning lay in exploiting neglected wells, forgotten passageways, and overlooked fissures, precipitating disaster from the unguarded. Today's cybercriminals demonstrate a similar acuity, eschewing direct attacks on an organisation's crown jewels and, instead, probing every unmonitored niche and shadowy recess within the corporate perimeter. Assets deemed negligible, systems considered deprecated, and endpoints assumed safe- these now constitute the new portals of compromise. For Chief Information Security Officers, Chief Information Officers, and the analyst community, the imperative is definite, attention must shift towards those assets long consigned to obscurity.

The Cybercrime Playbook - Reimagining Targets

The relentless advance of digital transformation, spanning cloud migration, integration of the Internet of Things, and an expanding universe of SaaS applications—has enabled both innovation and the

unchecked diffusion of organisational assets. This transformation, far from simplifying the defensive mandate, has proliferated an ecosystem of so-called 'shadow assets'-

- Cloud accounts established for development purposes yet never decommissioned,
- Legacy endpoint devices no longer maintained, yet still connected,
- Outmoded websites, orphaned domains, and user accounts surviving the offboarding of employees.

[Gartner's recent research](#) foregrounds this reality, highlighting that shadow and low-priority assets now underpin many initial breach vectors, frequently exploited as the attackers' first foothold. Such assets, though trivial to internal custodians, present open invitations to the adversary.

Incidents from the preceding two years reaffirm this paradigm shift. Sophisticated ransomware infiltrations, for instance, typically trace their origin to dormant endpoints or neglected third-party links. Even an unpatched cloud server, omitted from routine sweeps, has acted as the very entry point for data exfiltration at several multinational enterprises. These occurrences are ceasing to be statistical anomalies; they are fast becoming orthodoxy.

The Shortcomings of Traditional Risk Frameworks

Conventional approaches to vulnerability management—predicated on severity hierarchies and compliance tabulations, no longer suffice in this emerging risk landscape. Scheduled vulnerability scans may identify risks of high or critical importance, yet continue to disregard those assets which fall outside formal inventories.

This oversight can have dire consequences. Attackers frequently aggregate minor exposures, a forgotten virtual machine, an abandoned SaaS subscription, stale credentials, to construct consequential privilege escalation and lateral movement, breaches that static or periodic risk exercises are ill-equipped to predict.

Continuous Threat Exposure Management - A Standard Shift

Security leaders are responding by embracing Continuous Threat Exposure Management (CTEM) as the new essential. Differentiating itself from legacy paradigms, CTEM represents not merely a technological innovation, but an operational discipline, one that brings end-to-end visibility, validation, and business-relevant risk prioritisation to every corner of the digital estate.

Gartner defines the key requirements of an effective CTEM programme as follows-

- **Comprehensive scope:** Inclusion of all assets—legacy, cloud, shadow, and third-party.
- **Incessant discovery and profiling:** Constant detection to ensure new assets are swiftly catalogued and assessed.
- **Business impact prioritisation:** Assigning remediation urgency not by technical gravity alone, but by potential effects on business continuity and data integrity.

- **Active validation:** Employing attack simulation and red-teaming to discern the genuinely exploitable from the merely theoretical.
- **Operational integration:** Ensuring that findings translate into ongoing, real-world risk reduction—not simply endless identification cycles.

[Gartner](#) projects that by 2026, organizations implementing CTEM will be three times less likely to suffer a major breach than those relying on static, event-driven monitoring.

Analyst Perspective - Shifts in Budget, Prioritisation, and Regulation

CTEM's rise occurs in tandem with major movements in cyber risk investment and regulatory expectation:

- **Budget allocation:** Forrester's most recent benchmarks indicate that system defence activities—now encompassing automation, attack surface management, and advanced vulnerability mitigation—absorb 29% of total cybersecurity expenditure, surpassing both endpoint and cloud defence.
- **Risk-driven spending:** Funding increasingly aligns to exposure analytics, identity and access controls, and attack surface management, reflecting their actual operational impact, not merely technical classification.
- **Regulatory elevation:** Policymakers internationally, including the United States Securities and Exchange Commission and European Union authorities, are mandating risk-based vulnerability management, specifically auditing organisations' stewardship of so-called “immaterial” assets (Forrester).

The era of attempting universal remediation is at an end. The imperative now is to apply finite effort to only those exposures with genuine potential for exploitation.

The AI Catalyst - Detection to Defence at Scale

Artificial Intelligence is no longer a complementary element in enterprise defence—it is rapidly becoming foundational. As the scale and volatility of today's asset environment outpace manual capabilities, AI emerges as the only viable path from detection to defence at scale. It revolutionizes asset management through real-time monitoring and classification, enabling autonomous tools to identify and categorise assets regardless of their location or origin. It proactively detects and triages threats, using machine learning to uncover anomalies in low-priority or rarely accessed assets that traditional reviews often overlook. AI also orchestrates automated responses—from patch deployment to access revocation—delivering a speed and precision far beyond human reach. However, this advantage is not exclusive; generative AI is equally available to adversaries, powering large-scale reconnaissance, rapid exploit development, and automated attacks on overlooked assets. In this escalating AI arms race, defensive automation is no longer a strategic advantage—it is a fundamental necessity.

Case-in-Point - The Launchpad Role of 'Low-Priority' Assets

Contemporary breach investigations consistently highlight a pattern. Attackers have shifted from direct confrontation with hardened perimeters towards exploiting under-regarded auxiliary systems. In numerous publicised incidents:

- Orphaned cloud servers, excluded from standard patching cadences, have been paired with compromised credentials to facilitate lateral traversal,
- Unsecured test environments, overlooked after project close, have served as vectors for data extraction neither foreseen nor guarded by production controls,
- Legacy VPN gateways, retained for contingency purposes, have allowed multi-stage ransomware infiltration.

Such episodes are seldom isolated missteps. Successful attackers routinely exploit a chain of seemingly inconsequential exposures, demonstrating that security is only as robust as its least valued component.

Strategic Priorities for Security Leadership

For those entrusted with the enterprise's digital stewardship, the mandate has evolved to reflect the growing complexity of the threat landscape. Organizations must begin by exposing the unseen through exhaustive asset inventories, recognizing that any unmonitored asset poses a potential risk. This should be coupled with a shift from periodic audits to continuous discovery, ensuring real-time, integrated identification of assets and exposures. Prioritization must be aligned with business imperatives—technical severity alone is insufficient; systems deemed non-critical by IT may be essential to operations elsewhere. To enhance responsiveness, enterprises must institutionalize AI and automation, not just for detection, but for escalation management, root cause diagnosis, and autonomous remediation. Shared accountability must also be fostered, with DevOps, IT, and business stakeholders jointly responsible for securing both visible and shadow assets. Finally, it is essential to demonstrate regulatory diligence by aligning Continuous Threat Exposure Management (CTEM) with evolving compliance frameworks and delivering actionable insights to leadership and governance bodies.

Epilogue

Cybercriminals excel at uncovering fragile seams left unattended. The professional obligation of defenders is thus to erase any distinction between 'low-priority' and 'high-priority'—to recognise that every asset is consequential, and that exposure is dynamic rather than static. We must commit to vigilance in those once-overlooked domains, for only then shall we deny the adversary their clandestine passage.

About the Author

Pankit Desai is an entrepreneur and the Co-founder and CEO of Sequaretek, a cybersecurity, cloud security products and services company. He has been instrumental in growing the company into a leading provider of cybersecurity and cloud security solutions. Before starting Sequaretek, Pankit held various technology leadership and management roles in the IT industry across companies such as NTT Data, Intelligroup, and Wipro Technologies. He holds a degree in computer engineering and has a strong background in technology and entrepreneurship.

He can be reached at <https://www.linkedin.com/in/pankit>. Company URL- <https://sequaretek.com/>





Grappling with a Post-CVE World

Navigating Security Beyond Disclosure: Resilience, Response, and the Future of Cyber Defense

By Tod Beardsley, VP Security Research, runZero

When most people think of vulnerability management, they immediately think of the Common Vulnerabilities and Exposures (CVE™) program. For over a quarter century, CVE identifiers have become synonymous with tracking the enterprise's cybersecurity stance, forming a foundational pillar of security programs worldwide.

However, [earlier this year](#), this fundamental bedrock of cybersecurity was shaken when MITRE's National Security federally-funded research and development center (FFRDC) nearly lost the contract funding from the US Department of Homeland Security. A last hour intervention from the Cybersecurity and Infrastructure Security Agency (CISA) averted the worst-case scenario of shutting down CVE, but this crisis was a wake-up call for the cybersecurity industry.

While the CVE Program's continued operation remains critical to global cybersecurity efforts, and its closure would be a significant hit to tracking known vulnerabilities, we really need to come to terms with the fact that not all hacker tactics are described as CVEs. In fact, according to the [2025 Verizon DBIR](#), only about 20% of reported incidents can be traced to an exploited vulnerability for initial access.

What is required is a single source of truth when it comes to vulnerability and exposure management, and one that reflects the real-world risk landscape – not just CVEs.

The CVE crisis

CVEs give you a snapshot of enterprise assets, but they fail to provide a complete picture. They overlook critical issues like misconfigurations, segmentation flaws, and internally exposed assets, all flaws that attackers could exploit. Traditional tools fall short when it comes to asset discovery. Agent-based and credential-dependent solutions struggle to detect shadow IT, operational technology (OT) and IoT devices, all of which are increasingly common on today's attack surfaces, and difficult to monitor with traditional endpoint detection and response (EDR) and authenticated scans.

With under-resourcing a problem across the board, affecting not only the CVE program but also the [National Vulnerability Database](#) (NVD), an approach that isn't entirely CVE-centric is urgently needed.

Rising risk

This urgency is amplified by the complex nature of the modern corporate attack surface, which has become a tangled web of on-premises servers and desktops, remote working laptops and smartphones, public cloud containers, edge devices, and operational technology (OT). It is virtually impossible to maintain visibility and detect exposures with so many transient and dynamic assets and defenders are constantly left in the dark.

This is taking place in tandem with major changes to the threat landscape, which is becoming increasingly dangerous as actors grow more sophisticated and professional.

The cost

The consequences of cyberattacks are escalating and impossible to ignore. In the US alone, data compromises have reached a [near-record high](#), with almost 1.4 billion victims receiving notifications regarding a breach. Ransomware also remains a top concern, and recent research by Sophos indicates that [half of 3,400 responding IT professionals](#) paid ransomware operators in the first part of this year. The cost of the ransomware payment itself is just the tip of the iceberg, beyond this there is the business interruption, cost of missed sales and IT and legal costs.

When breaches stem from preventable exposures, organizations also risk facing regulatory penalties. Senior managers can be held personally accountable for instances of serious negligence and organisations can face huge fines and reputational damage if they can't give evidence which proves all assets are visible and secure.

Regaining the upper hand

The writing is on the wall: an over-reliance on CVEs and agent-based approaches won't keep you safe. So what else can you do to regain the upper hand?

Combining active scanning, passive discovery and API integrations is an effective method for gaining comprehensive visibility into both the internal and external attack services, including unknown and unmanaged assets like OT and IoT endpoints.

Once identified, the next step is to profile each asset in depth. This is when fingerprinting technology can play an integral part in extracting context-rich data. The more expansive the research is into what service a device uses, who the asset owner is, whether it's unpatched or misconfigured and what it's connected to, the more accurate the insight. This enables exposures that may otherwise remain an enigma to network defenders, to be understood.


Above all, solutions must be simple and data driven. That means consolidating capabilities into a single platform that has the capacity to deliver risk-based, prioritized alerts. Security teams are already overwhelmed by false positives, alert fatigue, and situational blindness, and what they need is to cut through the noise and see what exposures and vulnerabilities truly poses a threat.

About the Author

Tod Beardsley is the VP of Security Research at runZero. He "kicks assets and fakes frames." Prior to 2025, he was the Section Chief for the Vulnerability Response section for CSD/VM/VRC at [CISA](#), the Cybersecurity and Infrastructure Security Agency, part of the US government. He's also a founder and CNA point of contact for [AHA!](#). He spends much of his time involved in vulnerability research and coordinated vulnerability disclosure (CVD). He has over 30 years of hands-on security experience, stretching from in-band telephony switching to modern ICS/OT implementations. He has held IT ops, security, software engineering, and management positions in large organizations such as the Rapid7, 3Com, Dell, and Westinghouse, as both an offensive and defensive practitioner. Tod is a [CVE Board member](#), has authored several [research papers](#), and hosted Rapid7's *Security Nation* podcast with Jen Ellis. He is also a Travis County Election Judge in Texas, and is an internationally-tolerated [horror fiction expert](#).



Tod can be reached online at <https://www.linkedin.com/in/todb/> and at our company website <https://www.runzero.com/>



Guardians of the Factory: Defending Cyber-Physical Systems in Smart Manufacturing

By Omkar Bhalekar, Sr Network Engineer, Tesla

The modern industrial era blurs the lines between tangible machinery and interconnected computer systems. Step forward to the era of [Cyber-Physical Systems](#) (CPS), the smart integration of computation, networking, and physical systems. At the helm of this revolution is the [Internet of Manufacturing Things](#) (IoMT), an interconnected web of devices, sensors, and machines promising unparalleled automation, productivity, and responsiveness.

But with increased connectivity comes increased exposure. The same facilities that facilitate convenience of manufacturing processes bring new surfaces to breach [cyber-attacks](#). As factories become smart and integrated, it is now no longer an option but a necessity to secure these CPS and the IoMT ecosystem.

Understanding Cyber-Physical Systems in Manufacturing

Cyber-Physical Systems represent a convergence of the physical and virtual domains. In manufacturing, this equates to physical equipment, from industrial robot arms and conveyor belts to sensors and control systems, that is also computational, and networking enabled.

IIoT goes a step further by networking such devices via the internet or internal industrial networks such that they enable real-time data gathering, analytics, remote monitoring, and even autonomous decision-making.

Examples include:

- Intelligent sensors monitoring temperature, vibration, and wear
- Product-delivering [AGVs](#)
- Robot assembly lines with adaptive capabilities that adjust in real-time to production schedules
- [Digital twins](#) simulating and optimizing manufacturing processes

All this equipment makes productivity greater but also leaves vulnerabilities in one machine to impact the whole system.

Why Are CPS and IIoT Targets?

1. Critical Assets: Factory equipment is likely to produce mission-critical components in aerospace, automotive, and defense sectors. Supply chain interruption or IP theft has disastrous economic and geopolitics consequences.
2. Legacy Systems: Legacy equipment without security in their design are common in most factories. The devices may have obsolete software or no authentication whatsoever, making them a priority target.
3. Varied Attack Surface: CPS unites [IT and Operational Technology](#) (OT), which is typically managed by different teams with different security policies. It's hard to implement uniform security policies.
4. Increasing Connectivity: The demand for cloud integration, remote connections, and vendor relations expands the perimeter, presenting multiple points of attack to the adversary.

The Threat Landscape

Ransomware and Disruption

Ransomware now affects office networks no longer. In 2021, a global food processing behemoth was its production lines paralyzed by a ransomware attack for days, resulting in shockwaves through supply chains globally. Cyberattackers target manufacturing systems in order to be able to exert maximum pressure and demand huge ransoms.

Industrial Espionage

Intellectual property and trade secrets are sought to be stolen by nation-states and cybercriminals. Exposed IIoT devices can be used as covert spy gadgets or as vectors for long-term data exfiltration.

Manipulation of Physical Processes

Enemies can hack sensor feedback or command messages to cause machines to crash, produce defective products, or even become safety hazards. The Stuxnet worm was a pioneering case of such an attack, but today's attacks such as [FrostyGoop malware](#) are much more diversified and advanced.

Supply Chain Attacks

Vulnerabilities may be injected during the manufacturing or software update of IoMT devices themselves. These [supply chain vulnerabilities](#) are used by the attackers to gain access to industrial networks in general.

Strategies for Securing CPS and IoMT

1. Segmentation and Network Design

Isolate IT and OT networks wherever feasible, using firewalls and data diodes to restrict unauthorized access. Implement [micro segmentation](#) to restrict lateral movement between networks.

2. Zero Trust Principles

Adopt a [zero-trust approach](#) by treating all users and devices as unverified. Implement stringent IAM policies, multi-factor authentication, and real-time activity tracking.

3. Device and Firmware Security

Implement tamper-evident hardware and [secure boot protocols](#) in every device. Update firmware periodically and cryptographically sign the updates.

4. Behavior Analysis and Anomaly Detection

Implement AI-driven monitoring solutions that detect device behavior patterns and alert teams of possible anomalies that signal a possible breach or system failure in real-time.

5. Supply Chain Risk Management

Properly screen all your vendors, use robust security best practices for your vendor infrastructure, and beware of unauthorized firmware updates or unauthorized hardware revisions.

6. Incident Response and Recovery Planning

Design effective response procedures with definite roles and job profiles and perform recovery simulation on a regular basis to effectively manage breaks and continue business as usual when breaches take place.

The Human Factor

While technology is at the forefront, human capabilities are the foundation stones of CPS security. IT-OT cross-training, teamwork, and cybersecurity training are all essential ingredients in prevention and response to threat.

Looking Ahead: The Future of Secure Manufacturing

The future of CPS and IoMT is merciless, where factories will be smarter, speedier, and more responsive than ever. But it all depends on the security foundation in each layer, from device and chip architecture right through [sustainable manufacturing](#), network design, and operational procedures.


Upcoming and new technologies such as blockchain which enables transparency across secure edge computing and supply chains will continue to be the heroes in safeguarding manufacturing ecosystems.

Final Thoughts

At the crossroads of cyberspace and physical realm, the stakes have reached unprecedented heights never seen before. Cyber-Physical Systems are the backbone of industry manufacturing tomorrow yet must be safeguarded with caution, foresight, and ingenuity. Safeguarding the Internet of Manufacturing Things is not just about safeguarding equipment, it's about safeguarding industry's future

About the Author

Omkar Bhalekar is a Sr Network Engineer at Tesla Motors. He specializes in advanced networking in the field of smart and sustainable manufacturing technologies of electric Vehicles, battery technology, and Robotics following Industry 5.0 standards which focusses on environmental sustainability. With 7+ years' experience in this field and cybersecurity enthusiast specializing in Data center architecture, Manufacturing infrastructure, and Sustainable solutions, Omkar has extensive experience in designing and securing resilient industrial networks, building smart factories and AI data centers with scalable networks, Omkar avidly writes to simplify complex technical topics for engineers, researchers, and industry leaders. He is also the author of the Book: "Autonomous and Predictive Networks- The Future of Networking in the Age of AI"



Omkar can be reached online at LinkedIn: <https://www.linkedin.com/in/omkar-bhalekar>

or at his personal website <https://www.omkarbhalekar.com/>



How Digital Twins Enhance Grid Security

By Zac Amos, Features Editor, ReHack

Sophisticated cyber threats continuously threaten critical power infrastructure and are intent on causing widespread disruptions. As reliance on the electrical grid increases with the advancement of electric vehicles, data centers and smart technologies, more robust cybersecurity measures are necessary to keep pace with the modern world's interconnectedness. Digital twins are a next-generation solution to these demands and threats, enabling continuous monitoring, response and simulation, among other capabilities.

The Importance of Cybersecurity for the Power Grid

In 2024, [United States utilities experienced 1,162 cyberattacks](#) — a 70% increase from the year before. Although the attacks did not have much impact on the grid, the events raised the alarm about its susceptibility amid expansion and the integration of new technologies.

The electrical grid has several vulnerabilities, including its distribution systems. Transmission lines that carry electricity to customers are broadly outdated, and their [mechanisms enable remote access](#) to businesses, creating an opening for bad players to infiltrate and disrupt operations.

If the power grid fails due to a cyber attack, it will negatively affect governments, businesses and communities. Essential services — emergency response, defense and public safety — will halt, while adversaries might see an opportunity to threaten national security. Hospitals, transportation systems and wastewater treatment plants will also lose power, endangering public health.

The economy also has much to lose — such as data, supply chain stability and regulatory compliance — during widespread power disruptions. Studies show that just one hour of downtime [costs manufacturers \\$500,000 to \\$5 million](#), particularly in the pharmaceuticals, automotive, steel and chemical sectors.

Digital twins could offer much-needed protection. These tools can simulate the grid's behaviors so that engineers can better understand how it works. The digital twin uses real-time insights to generate an immersive environment, enabling simulations, strategic forecasting and improved decision-making.

Top Ways Digital Twins Enhance Grid Security

The digital twin market is rapidly expanding, helping organizations improve efficiencies and lower costs. These five capabilities are invaluable as the utilities sector strives for greater grid security.

1. Real-Time Threat Detection and Response

Digital twins can monitor grid operations in real time [to detect anomalies and deviations](#) from normal behaviors, flagging them for further investigation.

Integrating digital twins and machine learning into cybersecurity processes authorizes grid security teams to identify and address threats as they appear, allowing specialists to act quickly, isolate the problem and respond promptly.

2. Advanced Simulation of Cyber Attacks

Digital twins can simulate wide-ranging cyber attacks without compromising the original grid system. Utility companies can then use the simulation results to pinpoint vulnerabilities and assess the performance of security controls.

Modeling known and emerging threats provides crucial information about potential attack vectors and vulnerabilities, allowing power facilities [to ramp up their security postures](#) and best prepare for real incidents.

3. Predictive Maintenance and Risk Mitigation

Utilities can leverage artificial intelligence and machine learning in digital twins to predict grid failures and potential cyberattacks before they happen. Predictive analytics utilizes statistical algorithms that [determine the probability of future incidents](#) based on past data.

Data collection enables this tool to analyze vast amounts of information regarding operations and security. Utility companies can reduce unplanned downtimes and block exposure to vulnerabilities while boosting grid resilience and reliability.

4. Improved Compliance and Reporting

Digital twins streamline regulatory compliance by integrating automated collection, analysis and reporting. These technologies create comprehensive change tracking and documentation for industry and government standards, reducing human error and alleviating administrative work for staff.

The simplified process further enhances transparency and responsibility, allowing organizations to demonstrate their commitment to protecting grid infrastructure with robust cybersecurity measures.

5. Enhanced Collaboration Across Teams

Collaboration between information technology professionals, operational technicians and cybersecurity teams is critical for protecting the grid from cyber threats. Digital twins strengthen awareness and foster cross-functionality so everyone works together more effectively.

Several teams use a standard dataset and visualization software to identify potential issues and share their understanding. The integrated method provides a more proactive approach to grid cyber defense and collective accountability.

Challenges and Considerations

While digital twins offer enhanced cybersecurity benefits for the power grid, industry professionals must address several challenges. For example, digital twins must access, store and transfer [large quantities of sensitive data](#) that risk exposure if they are not properly managed and secured.


Ensuring systems are compatible with legacy grid infrastructure presents another hurdle. Digital twin technologies are also a costly investment, and maintenance expenses could be substantial. Smaller utility companies may have a difficult time covering the costs. Organizations of all sizes must weigh the pros and cons of deploying digital twins to maximize their cybersecurity efforts.

The Future of Grid Security with Digital Twins

Too much is at stake when it comes to safeguarding the power grid. It is time to evolve legacy cybersecurity practices to meet the complexity of modern cybercrime. Digital twin technologies present a viable solution with more sophisticated measures for protecting this critical infrastructure.

About the Author

Zac Amos is the Features Editor at ReHack, where he covers cybersecurity and the tech industry. For more of his content, follow him on [X \(Twitter\)](#) or [LinkedIn](#).





Breaking Traditional Encryption Protocols: Quantum Computing and the Future of Secure Communications

By Joe Guerra, M.Ed. CASP+, CCSP, FedITC, LLC

Introduction

Envision our digital world as a sprawling universe filled with glittering constellations of encrypted data, each packet of information orbiting in a precisely ordered, mathematically defined trajectory. Our job as cybersecurity professionals, especially those in Information System Security Officer (ISSO) and Information System Security Manager (ISSM) roles, is much like the work of cosmic gatekeepers, ensuring that these celestial bodies of data are protected from the chaos beyond.

For decades, classical cryptography has been the gravitational force holding this digital universe together. These systems rely on **classical bits**, which exist in only one of two binary states, a one or a zero, much

like a light switch being either on or off. That binary certainty has been the bedrock of our security. But on the horizon, there is a profound change, a disruptive, almost cosmic-level event in the form of **quantum computing**.

Unlike classical bits, **quantum bits** or **qubits** can exist in multiple states simultaneously thanks to principles known as **superposition** and **entanglement** (National Institute of Standards and Technology [NIST], 2023). This fundamentally alters the rules of the cryptographic game. The sheer computational power of quantum machines could soon dismantle encryption protocols that we once considered unbreakable. This isn't just an upgrade in processing speed, it's a rewriting of the universe's digital laws, demanding new strategies for survival.

The Fundamentals of Quantum Computing

To understand the threat, we must first journey into the “subatomic” realm of computation, where the boundaries of conventional logic blur, and probability takes the lead role.

Qubits are the basic units of quantum information. Unlike traditional bits, which can only store a 1 or a 0, qubits can store *both at the same time*—up until the moment they are measured. This is due to **superposition**, a phenomenon that allows multiple potential realities to coexist.

Imagine trying to find your way out of a massive maze. A classical computer would try one path at a time until it found the exit. A quantum computer, leveraging superposition, travels *all* paths at once and determines the solution in a fraction of the time.

But the real magic happens when you introduce **entanglement**—a phenomenon Albert Einstein famously called “spooky action at a distance.” Entangled qubits remain linked no matter how far apart they are, such that a change in one instantly affects the other. In computing terms, this allows for incredibly coordinated, parallel processing that vastly increases efficiency (Yuan et al., 2017).

The third key concept is **quantum parallelism**, which gives quantum computers their massive theoretical advantage. Because of superposition, they can work on an astronomical number of possibilities simultaneously, compressing centuries of classical computation into potentially seconds or minutes.

Some key concepts in short:

- **Qubit:** Infinitely more expressive than a bit, can be 0, 1, or both.
- **Superposition:** Multiple states exist in parallel until observed.
- **Entanglement:** Linked particles share the same state instantaneously.
- **Quantum Parallelism:** Multiple possibilities tested simultaneously.

Current research from industry leaders such as IBM, Google, and IonQ is pushing towards what is known as **quantum advantage**—the tipping point where a quantum system can perform a calculation that is impractical for even the largest classical supercomputer (NIST, 2023). Once *fault-tolerant* quantum computing is achieved, formerly invincible encryption methods may fall rapidly.

Technical Threats: Shor's and Grover's Algorithms

The threat is best understood through two quantum algorithms—**Shor's** and **Grover's**, which can be seen as the “unlocking formulas” of the quantum era.

Shor's Algorithm (1994) revolutionized theoretical cryptanalysis by demonstrating that a sufficiently powerful quantum machine could factor large integers and solve discrete logarithms exponentially faster than classical methods. This is devastating news for RSA and ECC encryption, which rely on these mathematical problems being computationally infeasible (NIST, 2023).

As an example, RSA-2048, considered secure because factoring such large numbers would take billions of years for classical hardware—could theoretically be broken in minutes with a large enough quantum computer. The implications for banking, e-commerce, and secure government communications are staggering.

Grover's Algorithm, while not as directly destructive to public-key crypto, still poses a significant threat. It offers a quadratic speed-up for brute-force searches, effectively halving the security level of symmetric-key algorithms and hash functions like SHA-2 and SHA-3 (NSA, 2022). This means that to maintain equivalent security, organizations will need to double their key sizes.

Algorithm	Target	Quantum Advantage	Security Impact
Shor's Algorithm	RSA & ECC encryption	Exponential speed-up factoring	Breaks RSA-2048/ECC-256 in hours
Grover's Algorithm	Hash functions (SHA-2, SHA-3)	Quadratic acceleration brute-force	Reduces effective bit strength by ~50%

The Real and Hypothetical Impact

The quantum threat extends across multiple domains:

Finance: Online banking, stock exchanges, and payment processors rely heavily on RSA and ECC. A quantum-enabled breach could allow not just theft but digital fabrication of valid transactions—undermining trust in the financial system.

Government and Defense: Public Key Infrastructure (PKI) secures sensitive inter-agency communication. Under the *harvest-now, decrypt-later* threat model, adversaries can store encrypted communications today with the intention of decrypting them once powerful quantum technology is available (NSA, 2022).

Healthcare: Electronic Health Records (EHRs) and PHI are often protected by ECC-based security. Quantum compromise would expose patient identities, diagnoses, and histories, ripe for exploitation or blackmail.

Intellectual Property: Corporations could have trade secrets and patent applications stolen before filing, allowing competitors to replicate and deploy them first.

Critical Infrastructure: Power grids and water systems rely on PKI for SCADA system commands. A quantum-empowered attacker could feasibly disrupt these at scale, causing cascading failures in essential services.

Sector	Vulnerable Encryption	Quantum Threat Impact	Example Incident Scenario
Finance	RSA/ECC	Theft, forgery, payment compromise	Fake SWIFT messages; intercepted transfers
Government	PKI	Espionage; secret exposure	Decrypting stored intelligence after Q-Day
Healthcare	ECC/PKI	Identity theft, privacy breach	Nationwide hospital records breach
IP/Corporate	PKI	Competitive espionage loss,	Patent data stolen pre-filing
Infrastructure	PKI-secured SCADA	Physical sabotage, outages	Quantum-enabled blackout

Preparing for the Quantum Era

Defending against quantum threats means starting now:

1. **Inventory & Classification** – Map all instances of vulnerable encryption across your organization. Identify high-risk systems.
2. **Post-Quantum Migration** – Implement NIST-selected algorithms like **Kyber** and **Dilithium**, designed to be resistant to quantum attacks (NIST, 2023).

- 3. **Crypto-Agility** – Architect networks to pivot quickly to new cryptographic standards, reducing transition time during protocol updates.
- 4. **Quantum Key Distribution (QKD)** – Leverage quantum physics itself to create unbreakable key exchanges (Yuan et al., 2017).
- 5. **Layered Encryption** – Combine robust symmetric systems like AES-256 with quantum-safe public keys, ensuring multiple layers of protection.
- 6. **Training & Awareness** – Prepare ISSM/ISSO teams with regular workshops on quantum threats, migration strategies, and evolving NIST guidance.

Global Case Studies

Quantum readiness is no longer hypothetical, nations and corporations are already experimenting with quantum-secure communications. China’s **Beijing–Shanghai QKD backbone**, a 2,000 km fiber optic network, demonstrates successful long-distance deployment of quantum encryption at a national scale (Yuan et al., 2017).

Major banks like J.P. Morgan and HSBC are piloting quantum-safe cryptography in their transaction networks, while tech giants explore hybrid solutions blending post-quantum algorithms with classical methods (NIST, 2023).

These initiatives show that preparedness is a strategic advantage. Those who delay risk scrambling post–Q-Day, with regulatory penalties, operational disruptions, and catastrophic trust erosion.

Feasibility, Limits, and Timelines

While still developing, quantum computing is advancing along a clear trajectory:

Stage	Capability	Estimated Arrival	Security Implication
Research	Factor small keys (PoC)	2001–present	No immediate threat
Pre-Advantage	100–500 noisy qubits	Today	Classical crypto holds ... for now
Quantum Advantage	Thousands of fault-tolerant qubits	~2030	RSA/ECC-breaking feasible

Stage	Capability	Estimated Arrival	Security Implication
Post-Q-Day	Mass decryption capability	5–15 years	Global crypto overhaul required

The engineering challenges, qubit stability, error correction, scalability, are being addressed at an accelerating pace. The likely window before Q-Day is narrow enough that critical sectors must act now (NIST, 2023; NSA, 2022).

Conclusion

Quantum computing represents both a marvel of physics and an existential risk to our digital civilization. Where once cryptography was the immutable law of cyberspace, the peculiar rules of quantum mechanics could soon overturn those laws entirely. State and criminal actors are already preparing—harvesting encrypted data for future decryption.

To survive this deflection point in cybersecurity history, ISSM and ISSO professionals must take leadership in quantifying cryptographic exposure, migrating to quantum-resistant systems, implementing crypto-agility, and embedding new defense concepts like QKD.

In the words of Michio Kaku, we are about to enter an age where “the impossible becomes inevitable.” For cybersecurity, this means the future will belong to those who prepare, not after the quantum wave hits, but before it crests on the horizon.

References

National Institute of Standards and Technology. (2023). *Post-quantum cryptography project*. <https://csrc.nist.gov/projects/post-quantum-cryptography>

National Security Agency. (2022). *Quantum computing and cryptography: Key policy issues*. <https://media.defense.gov>

Yuan, Z., Chen, Y., Zhao, Y., & Zhang, Q. (2017). The Beijing–Shanghai quantum communication backbone: Science, technology and innovation. *National Science Review*, 4(2), 193–200. <https://doi.org/10.1093/nsr/nwx004>

About the Author

Joe Guerra, M.Ed., CASP+, CCSP, RMF ISSO/ISSM Instructor, FedITC, LLC. San Antonio, Texas (Lackland AFB)

He is an experienced computer science and cybersecurity educator with over 20 years of expertise. He spent 12 years teaching science, Information Technology, and Computer Science at the high school level, shaping young minds and inspiring the next generation of technology professionals. His deep knowledge and passion for the field paved the way to higher education. Joe holds a Master's degree in Information Systems Security and Instructional Technology, and is certified in CompTIA Network+, Security+, CySA+, and CASP+, as well as CCSP by ISC2.



For the past 10 years, Joe has been an esteemed adjunct instructor at ECPI University, the University of the Incarnate Word, and Hallmark University. He has taught a wide range of courses, including Security Assessment and Testing, Identity and Access Management, Linux operating systems, and programming languages such as Java, C, Python, C#, and PowerShell. His diverse skills also encompass networking, cybersecurity, Cisco systems, Hacking and Countermeasures, and Secure Software Design.

A highlight of Jose's career was his 2019–2023 role teaching Air Force cyber capability developers, where he focused on developing offensive and defensive software tools, making significant contributions to cybersecurity warfare and national defense.

In addition to his technical teaching, Joe specializes in training cyber leadership personnel, including Information System Security Officers (ISSOs) and Information System Security Managers (ISSMs), in the Risk Management Framework (RMF) process. He equips these cyber professionals with the knowledge and practical skills required to navigate complex regulatory environments, ensure compliance with federal standards, and implement robust security controls. Joe's instruction emphasizes real-world application of RMF, fostering a deep understanding of risk assessment, security authorization, and continuous monitoring. His approach prepares ISSOs and ISSMs to lead cybersecurity initiatives, manage enterprise risk, and uphold the highest standards of information assurance within their organizations.

Joe's dedication to education, hands-on expertise, and leadership in both technical and managerial aspects of cybersecurity make him a trusted mentor and resource for students and professionals alike. Joe can be reached online at jguerra@Feditc.com, joe.guerra@afcsu.org, mrjoe Guerra@outlook.com and at our company website <https://feditc.com/>



IGA Is a Journey, not a Destination – Plan Accordingly

Success with IGA requires a shift in mindset from project to process, from quick fixes to strategic alignment.

By Niels Fenger, Advisory Practice Director, Omada

Done properly, Identity Governance and Administration (IGA) brings huge benefits to an enterprise. IGA enables IT and security teams to continuously monitor, manage and administer access rights to a constantly evolving set of resources. It is a critical tool for any organization looking to reduce its risk, meet compliance and enable its workforce, especially as digital identities proliferate.

However, many organizations still struggle with their IGA investments. When it comes to IGA implementations, one of the biggest problems continues to be the way IGA is treated within the organization. In other words, the source of failures and foibles often stems from approaching IGA as a one-time event. The reality is that IGA is a journey, not a destination, if you want it to succeed.

Challenges with a one-and-done mindset

The core issue with treating IGA as a project is that projects have defined start and end points. That's not how IGA works. It's an ongoing process that needs to evolve alongside the organization.

Every day, employees join, change roles or leave, and each event requires updates to access rights. Meanwhile, most organizations retire or replace roughly 10% of their applications annually, meaning the landscape of systems and access requirements is always shifting. What's secure and appropriate today might be a compliance risk tomorrow.

Access is not a "set it and forget it" task. Roles change. Responsibilities shift. Business units restructure. All of this requires that access privileges be continuously reviewed, adjusted, and revoked as needed. And it must all be done in a documented, auditable way to meet the growing demands of regulatory frameworks and internal controls.

For IGA to be effective, it must be treated as a core operational function — one that is maintained by a dedicated IGA operations team, not run as a time-limited initiative.

Why IGA is often treated like a project

It often comes back to how IGA is sold or marketed. It's often presented as the solution to address a technical pain point. For instance, "How can we report that we manage all of our Sarbanes-Oxley (SOX) applications?" Or "How do we get approvals on all our access assignments?" Then, perhaps, someone starts researching the need and determines that an IGA solution can solve a particular problem.

What happens here is the IT team or whoever else is leading the charge will purchase the solution and think only of it as being meant to solve the specific problem — like the ability to report on SOX classified access. But IGA is so much more than an expensive report generator or an approval engine.

How to move from destination to journey

Ideally, rather than purchasing an IGA solution with only a specific problem in mind, the better approach is to look holistically at your organization's IT strategy and consider how IGA could support it overall. This helps determine priorities and, ultimately, helps align IGA to those priorities. IGA can map into ISO27001, NIST, NIS2, DORA, SOX, PCI, GxP and all the other security frameworks and industry-specific cybersecurity regulatory requirements.

Defined ownership is also key. When organizations struggle with IGA, another overarching trend is that there's a discrepancy over who is in charge of the IT. There is often confusion about who is ultimately responsible for IGA. To succeed, someone within the organization must be formally accountable for the IGA domain.

Once ownership is established, the IGA operations team can be empowered to execute the solutions that support the broader business, IT or IGA strategy. Adhering to core IGA principles is essential for long-

term success. Relying on quick fixes or basing governance on untrustworthy master data can lead to serious issues. Without a long-term perspective, organizations risk locking themselves into rigid, unmanageable solutions.

For example, if data classification is delegated to individual system owners without a unifying framework, any hope of consistent, organization-wide reporting and access management quickly dissolves. Classifications must be defined globally, supported by a solid business case. Once this foundation is in place, the IGA operations team can implement the classification scheme and ensure that relevant owners are responsible for maintaining accurate classifications for their systems.

Best practices for IGA success

If you're unsure how IGA fits into your overall IT or security strategy, now is the time to clarify its role. If your IGA strategy is clear but the guiding principles for it are not, it's time to define those principles. If your principles are well-established but every initiative is still being run through ad hoc project teams due to the absence of a dedicated IGA operations team, then it's time to build one.

No matter where you are on your IGA journey, it's crucial to partner with someone who has a proven track record at the level you're aiming to improve. While software matters, successful IGA implementation is only about 20% tooling and technique and 80% organizational alignment and execution.

Commit to the IGA journey

Treating IGA as a one-off project is one of the most common — and costly — missteps organizations make. It is not something you can "complete"; it's a living, breathing part of your security and IT ecosystem. As identities, roles, systems and regulations continuously evolve, so must your IGA program.

Success with IGA requires a shift in mindset from project to process, from quick fixes to strategic alignment. That means building the right foundation: clear ownership, strong operational capabilities, well-defined principles and the organizational will to treat IGA as a continuous discipline.

When done right, IGA enables far more than compliance checkboxes. It empowers organizations to move with confidence, meet ever-changing regulatory demands and provide secure, appropriate access — every day, for everyone. That's not a destination you arrive at. It's a journey you commit to.

About the Author

Niels Fenger is the Advisory Practice Director at Omada, where he helps customers shape their Identity Governance and Administration (IGA) strategies and apply best practices. With over 25 years of experience in IT, Niels has led projects and advised at every stage of the process — from defining requirements and building business cases to planning workflows and working with executive teams. His global experience has given him valuable insight into a wide range of business cultures and practices. Niels can be reached online on [LinkedIn](#) and at <https://omadaidentity.com>.





The Perimeter Is Not Enough

5 Steps to Mitigate Risk in a Zero-trust Environment

By Parker Pearson, Chief Strategy Officer, Donoma Software



Kees Streefkerk, Unsplash

Perimeter-based solutions always have vulnerabilities, and new ones are popping up every day through system and vendor updates. It's a moving target, and it's quickly becoming impossible to keep pace with the rate at which these vulnerabilities are being exploited. So, the question becomes, what can CISOs, and CIOs do to better protect their data?

In the field of cyber defense, there has always been a great deal of emphasis on defending the network perimeter – so much so, in fact, that many organizations still focus their entire cybersecurity strategy on perimeter-based solutions.

Beyond the blurring of physical boundaries, there are many other reasons why defending the perimeter is no longer enough. The entire concept of securing the perimeter is inherently reactive, and the proliferation of AI-accelerated threats has created a situation in which it's difficult (if not impossible) for defenders to keep up with the speed and volume of automated attacks coming in from the outside. Although software patching remains critical, each patch is written and distributed only *after* a vulnerability has been exploited. By then, the damage has already been done. It's like playing a perpetual game of Whack-a-Mole in the midst of a 24/7 Urgent Care Center.

But perhaps the most urgent reason to move beyond a perimeter-only approach is the reality that – even now, as you read this article – hackers are likely already silently inside your IT systems.

External vs. Internal Threats

Cyber threats can be separated into two main categories: External attacks and internal attacks. Unfortunately, for most organizations, their IT networks are still extremely vulnerable to both types.

- **External attacks** are perpetuated by criminal persons, organizations or even nation state adversaries who find highly creative ways to gain access to systems so they can steal or ransom data. External attacks can take many different forms, but some of the most common include malware, phishing, ransomware, Denial-of-Service (DOS) and Man-in-the-Middle (MITM) attacks. In larger organizations, we also see SQL Injection, Zero-day exploits, and Spoofing attacks. In the first five months of 2025, [more than 22,000 new Common Vulnerabilities and Exposures \(CVE\) records](#) were received by the National Institute of Standards and Technology (NIST), with a backlog of nearly 25,000 more reportedly awaiting analysis. In 2024, over 40,000 new CVEs were published, up by more than 37% from 2023. These vulnerabilities are being exploited and weaponized faster every year on a massive scale.
- **Insider Threats** happen when individuals with credentialed access misuse their privilege – either intentionally or unintentionally – to harm the organization. These are perhaps one of the most concerning types of threat. [The Cybersecurity Insiders 2024 Report](#) indicates that 83% of organizations reported insider attacks in 2024, with 51% experiencing six or more attacks in the past year. In many cases, these are socially engineered, either by bribing existing employees or intentionally placing individuals as employees/contractors within an organization, just to gain access to data. Internal administrators often have inappropriate or unnecessary access to data that would normally be off limits to someone at their level of responsibility. Another form of insider threat is a Supply Chain attack, which targets third-party vendors or partners to compromise their products or services, which are then used to attack the main organization.

It's important that everyone in the organization be educated about how to protect against both types of attacks. Data is an asset, and its protection can no longer just be delegated to IT without oversight and understanding.

The Threat is Already Inside Your Network

Most organizational leaders want to believe that an internal threat isn't likely. After all, who wants to work with people who would steal? Sadly, most organizations learn the hard way, because they can't imagine that a data breach will happen to them. In fact, many organizations still view a data breach as an unlikely scenario, with odds similar to being hit by a tornado or a fire. So many assume that if they are not a household name with millions of customers, they will be poor targets. Still others, believe they are adequately prepared but never run simulated exercises to check. But the odds are that much higher – [nearly 1/3 of organizations will get hit by a data breach this year](#).

So why are so many breaches still occurring? The fact is, most people are working off some erroneous assumptions, particularly around their use of data encryption. While many security personnel proudly announce that all their data is encrypted in transit and at rest, what no one has been acknowledging is the dirty secret that *once systems are in use, all that protection goes away*. Most organizations never shut down their applications (even after hours, if they do not run continuously) so their data is *always* vulnerable, to anyone (known and unknown) inside the perimeter.

The [recent Coinbase data breach](#) is a great example. The breach did not result from a technical vulnerability in its systems, but rather was perpetuated from within, by support staff who abused their legitimate access to steal the data in return for relatively modest bribes. The breach compromised the sensitive personally identifiable information (PII) of almost 70,000 users, along with account-related information such as balance snapshots and transaction histories. This unauthorized activity happened over the course of almost 6 months before being discovered. As a result, Coinbase is facing at least six class action lawsuits alleging that Coinbase failed to implement and maintain adequate security protocols, exposing users to serious risks. In response to the breach, Coinbase has refused to pay the \$20 million US ransom demand and instead offered a \$20 million reward for information leading to the identification and prosecution of the attackers. The company estimates that the incident could cost between \$180 million and \$400 million, accounting for remediation efforts and reimbursements to affected users.

In another example, [Capital One experienced an enormous data breach](#) in 2019, due to a misconfiguration of their cloud infrastructure (specifically a misconfigured Web Application Firewall). This vulnerability was then exploited to access sensitive data from over 100 million customers, including credit scores and banking details. Unfortunately, this is not an isolated incident. System administrators often have too much access to organizational data compounded by them escalating their access privileges without the knowledge of management.

In another notable instance, in 2024, a hacker broke into AT&T's cloud storage provider, Snowflake, and accessed call and text records for almost all their 109 million US customers. Although AT&T claims that no names were attached to the stolen data, the breach [led to multiple class action lawsuits](#) were just recently settled for \$177 million US.

The reality is that we now must operate in a Zero-trust environment; but no organization can achieve full data privacy protection unless they also protect their data at its source. This means securing data via continuous encryption not just at rest or in transit. This is an open gap that needs to be addressed in every organization.

5 Steps to Mitigate Risk and Put the Brakes on Data Breaches

Let's face it: if perimeter-based security solutions were enough, we would not have daily data breaches. The problem is that perimeter-based solutions always have vulnerabilities, and new ones are popping up every day through system and vendor updates. It's a moving target, and it's quickly becoming impossible to keep pace with the rate at which these vulnerabilities are being exploited.

So, the question becomes, what can CISOs, and CIOs do to better protect their data? Here are five steps to take immediately:


1. **Create and Regularly Test Incident Response Plans.** CISOs and other technical leaders must collaborate with executive leadership to simulate real-world scenarios to test response capabilities, updated regularly to address evolving threats. Having a clear tested plan speeds response time and minimizes damage when attacks occur. Develop detailed incident response procedures, assign specific roles to team members, conduct regular tabletop exercises, and ensure all employees are trained to recognize and report suspicious activities quickly.
2. **Take a closer look at your data access policies.** In most organizations, the IT staff is spread thin, and administrators are expected to have expertise in a highly complex ecosystem of hardware and software. The high stress workload of these teams often result in corner cutting, such as unchecked access to all digital systems in the hands of an understaffed team who may not have the correct skills, either technical or operational to have such unfettered access. The IT function is often managed by the CFO or the COO, whose areas of expertise preclude them from understanding what their IT staff does on a daily basis and the many operational risks and threat vectors lurking in an area so many do not understand. Employees should be granted access only to what they need to do their jobs.
3. **Implement continuous monitoring.** Up until recently, it's been relatively easy for people to hide nefarious activities within the noise of systemic network checks. But now, with AI being woven into the fabric of the network, it's becoming easier to detect anomalous access and behavior patterns.
4. **Embrace Multi-Factor Authentication (MFA).** Weak authentication mechanisms or poorly managed access controls can lead to unauthorized access to critical systems and data. MFA provides a foundational security layer that can dramatically reduce the risk of account compromise. On its own it is not enough, but there are still many organizations who have yet to implement this first step to good data protection.
5. **Embrace Privacy Enhancing Technologies for continuous encryption.** Data encryption is a fundamental element of every viable enterprise cybersecurity strategy, but many executives are operating under the misconception that their data is safe because they've implemented something billed as "end-to-end" encryption. Unfortunately, most of these solutions only encrypt data while it is at rest or in transit. Once the application is in use, the encryption at rest ceases to be of benefit

and the data is exposed as clear text. Only continuous (homomorphic) encryption now available in new Privacy Enhancing Technologies (PETs) continuously protect data. By eradicating data loss, PETs can put an end to data breaches, data leaks, and the damage they cause. In doing so, it can eliminate the damage to brand reputation and ensuing litigation associated with data breaches, as well as regulatory reporting requirements. It not only helps organizations protect their intellectual property, confidential information and trade secrets from competitors, foreign adversaries and criminals – it preserves their brand value and bottom line.


In conclusion, perimeter-based cybersecurity solutions are like a moat around a castle. They may deter and defend, but they do not stop hackers from getting into networks or people with legitimate credentials from stealing data, plug security holes that leak data, or eliminate mistakes and misconfigurations of settings. Many security breaches occur because employees are unaware of best practices in cybersecurity or are careless in following them. Perhaps the single most important cybersecurity best practice for businesses to follow today is to invest in a multi-layered defense strategy, also known as “Defense in Depth.” This layering should include not just perimeter and Zero-Trust strategies, but also secure data down at its source via continuous encryption now becoming commercially available. This approach minimizes the impact of breaches by ensuring that as security controls are compromised, the data remains protected.

About the Author

As Chief Strategy Officer at Donoma Software, Parker Pearson leads the company's mission to revolutionize enterprise data privacy protection. With over twenty years of experience in technology innovation and entrepreneurship, she helps turn cutting-edge ideas into real-world solutions that solve complex business problems. At Donoma, she has helped power the company from startup to industry disruptor with their innovations in continuous encryption and privacy enhancing technology. Her expertise lies in making innovation understandable and aligned to meet evolving client challenges. She consistently seeks the next challenge to bring new ideas to life and transform them into commercially available solutions organizations need. What drives her most is the belief that curiosity, creativity and imagination are the true source of innovation. The rest is just ones and zeros.



Parker can be reached online at <https://www.linkedin.com/in/ppearson/> through the Donoma Software website: <https://www.donomasoftware.com/>.



Passwordless, but Not Riskless

Passwordless Authentication

By Sudhakar Tiwari, Principal Solutions Architect, Zurich

Modern enterprises are racing toward passwordless authentication, especially passkeys built on FIDO2/WebAuthn, to end the pain of passwords: fewer resets, far less reuse, and strong resistance to classic phishing. It's a worthy goal. But as any seasoned CISO will tell you, no authentication scheme is a magic shield. Passwordless dramatically shifts risk it doesn't eliminate it. Done right, it raises the security bar and improves UX. Done casually, it opens new blind spots. This piece explains why passwordless \neq panacea, highlights the new attack surfaces, and offers practical safeguards so you can reap the benefits without inheriting avoidable risk.

The Allure and the Reality of Passwordless

Passwords remain a top cause of breaches; attackers monetize guessable, reused, and phished secrets at scale. Passkeys replace shared secrets with public-key cryptography: a device-bound private key stays

local; the server stores only the public key. Binding to the correct origin means a look-alike site cannot complete the cryptographic challenge. No password database to steal, no credential stuffing, and vastly reduced phishing.

That promise is real but incomplete. Passwordless changes your threat model:

- You shrink the “remote phishing” surface.
- You increase dependence on endpoint integrity, device custody, and account recovery flows.
- You introduce cloud sync and enrollment risks that didn’t exist in the same way before.

Smart programs acknowledge those trade offs and engineer for them.

New Attack Surfaces in a Passwordless World

1) Device Theft & Token Loss

In passwordless deployments, the device is the credential. If a phone or hardware key is lost or stolen—and user verification (PIN/biometric) is weak or absent—an attacker might authenticate. Keys can be revoked, but only if enrollment and recovery are well designed and monitored. Unlike a password, you can’t “change a fingerprint”; you must revoke authenticators quickly and cleanly.

Mitigate: enforce device unlock (PIN/biometric), short screen-lock timers, disk encryption, and rapid revocation paths for lost authenticators. Issue backup keys kept offline in a safe place for business continuity.

2) Biometric Spoofing & Sensor Failures

Biometric systems can be spoofed or misread. Public demos have shown that even mainstream face or fingerprint systems can fail under specific conditions (many later patched). Biometrics are powerful, but fallible—and non-rotatable.

Mitigate: prefer on-device biometrics protected by secure enclaves; require liveness/anti-spoofing settings; combine with user-verification PIN for critical actions; and plan for fallback that isn’t easily phishable.

3) Malware & Endpoint Compromise

Passkeys live on endpoints. If an attacker gains OS-level control, they may hijack sessions, abuse on-device APIs, or tamper with UI to trick approvals.

Mitigate: treat passkey devices as protected assets: EDR/EPP, timely patching, app attestation where supported, least-privilege, and browser hardening. Monitor for anomalous authenticator use.

4) Cloud Account Exploitation

To enable cross-device use, many ecosystems optionally sync passkeys via cloud. If an attacker takes over that cloud account (weak recovery, SIM swap, social engineering), they may gain leverage over synced authenticators.

Mitigate: require strong, phishing-resistant MFA on the user's Apple/Google/Microsoft accounts; restrict or opt out of cloud sync for high-risk roles; scrutinize account-recovery changes.

5) Phishable Backups & Recovery

The weakest link is often a “convenience” fallback—SMS codes, email links, or helpdesk resets used when a user lacks their device. Attackers simply pivot to those flows, defeating the point of passkeys.

Mitigate: make every path (enrollment, recovery, device replacement, new-device sign-in) as phishing-resistant as the primary. Require supervisor/ID proofing for resets, enforce out-of-band confirmations, and disable SMS/email for privileged accounts.

6) Social Engineering & Coercion

No password doesn't mean no people. Attackers still push users to approve prompts, scan rogue QR codes, or enroll attacker-owned keys. In some cases, they add their own authenticator post-compromise to maintain persistence.

Mitigate: user education; approval fatigue protections; out-of-band verification for new authenticator enrollment; and alerts on enrollment events—especially for admins.

Passwordless vs. Phishing-Resistant: Know the Difference

- **Passwordless** = no memorized password (could be biometrics, magic links, OTPs, passkeys).
- **Phishing-resistant** = the method does not expose a reusable secret and won't authenticate the wrong origin (e.g., FIDO2/WebAuthn passkeys, smart cards, certificate-based auth).

Some passwordless methods (SMS, email links, OTPs) **are not** phishing-resistant. They may be useful transitional tools, but they're **not end-state security**. Aim for **passkeys** where feasible, and ensure **all auxiliary flows** meet the same bar.

Field Lessons (What Real Programs Learn)

- **Biometrics need hardening and patch discipline. Research keeps finding edge cases; vendors patch; you must stay current.**
- **Cross-device flows can be abused** if users aren't trained to recognize legitimate prompts/QRs. Proximity checks and trusted out-of-band confirmations help.

- **Attacker persistence via enrolled authenticators** is real. Treat **authenticator enrollment** like a sensitive transaction: log, alert, and verify changes.
- **Usability determines adoption.** Pilots stall without training, clear guidance, and thoughtful support for shared devices, gloved users, or legacy systems. Expect hybrid patterns in OT/clinical environments.

Best Practices: How to Deploy Passwordless Securely

Choose Truly Phishing-Resistant Methods

- Prefer **FIDO2/WebAuthn passkeys** and **hardware security keys** over OTPs/links.
- If using phone authenticators, use **platform passkeys** or **app-based** approvals with strong device attestation—not SMS.

Lock Down Devices & Accounts

- Enforce **PIN/biometric unlock**, **disk encryption**, **OS/browser patch SLAs**, and **EDR** on endpoints that hold passkeys.
- Protect consumer cloud accounts (that may sync passkeys) with **strong MFA** and **restricted recovery**; consider **no-sync** for high-risk personnel.

Harden Fallback & Recovery

- Treat **helpdesk, recovery, and enrollment** as **Tier-0** flows:
 - Require **phishing-resistant verification** (e.g., in-person, SSO-based proofing, secure video ID, or supervisor approval).
 - Disable SMS/email for **admins** and other critical roles.
 - Issue **backup hardware keys** and document secure custody.

Instrument & Monitor

- Log and alert on:
 - **New authenticator registrations**
 - **Repeated user-verification failures**
 - **Unusual geo/ASN patterns**
 - **Use of backup routes**
- Add **behavioral analytics**: deviations from typical device/origin/time can flag compromise.

Phase the Rollout

- Start with low-risk apps or pilot groups.
- Gather UX feedback; fix friction before expanding.
- Provide just-in-time education: what a real prompt looks like, why never to read a code to anyone, how to report a lost device fast.

Keep a Resilient Break-Glass Plan

- Define how to restore access if multiple authenticators are lost or a biometric fails:
 - **Break-glass accounts** with strict controls
 - **Emergency recovery** with multi-party approval
 - Rapid **revocation** + **re-enrollment** workflows
- Test this plan like you test DR.

Governance & Controls (for IAM Teams)

- Set **policy baselines**: which groups must use passkeys; which fallbacks are allowed; rotation/attestation requirements; logging minimums.
- Treat **authenticator enrollment** as a **high-risk change** requiring approvals and alerts.
- Map passkey posture into **identity risk scoring** and **access reviews**.

Leadership Playbook: Guidance for CISOs & Identity Architects

1. **Frame it as a journey.** Passwordless is not a switch; it's a **program** involving IAM, endpoint, HR (joiners/movers/leavers), legal/compliance, and helpdesk. Set **realistic milestones** and success metrics (adoption %, reduction in password resets, phishing incident rates).
2. **Run a risk assessment up front.** Answer:
 - What if a device is stolen?
 - How do we verify new device enrollment?
 - Which recovery flows are acceptable for which roles?
 - How do we monitor and revoke authenticators at scale?
3. **Use adoption to upgrade culture.** If users no longer need passwords, they can handle tapping a key. Provide **training**, normalize reporting lost devices quickly, and celebrate **security positive behavior**.

4. **Update incident response.** Add playbooks for:

- Revoking authenticators en masse
- Remote wipe / MDM workflows
- Containing abused recovery flows
- Rapid re-issue of keys to critical users

5. **Track the ecosystem.** Passkey standards and platform support evolve (proximity checks, attestation, anti-spoofing). Keep architecture reviews continuous, not one-and-done

Conclusion: Embrace the Future—Eyes Open

Passwordless especially passkeys is a major leap forward. It slashes entire classes of attacks and simplifies life for users and admins alike. But passwordless is not riskless. The threat frontier moves: to devices, enrollment and recovery, cloud sync, and human manipulation. Organizations that succeed won't just "turn on passkeys"; they'll engineer the whole lifecycle authenticator issuance, user verification, recovery, monitoring, and incident response with the same rigor once reserved for passwords.


Adopt the technology, absolutely. Pair it with policy, process, and monitoring, and you'll earn both the phishing resistance you want and the resilience you need.

About the Author

Sudhakar Tiwari is a Senior Cybersecurity Architect and IAM Strategist with over 18 years of experience in cybersecurity and Identity & Access Management.

In his current role, Sudhakar leads IAM strategy and solutions for global businesses, focusing on enhancing security while improving user experience. Designing secure access frameworks across Fortune 100 enterprises. His work focuses on application security, adaptive identity, and building secure-by-design architecture for critical industries. He is a Senior IEEE member, along with CISA and CISM security certifications

Sudhakar can be reached online at Sudhakar.tiwari2@gmail.com





Privileged Access in the Age of Ransomware-as-a-Service (RaaS)

Ransomware Has Become a Service Economy

By Sandeep Dommari, Principal Architect, Ping Identity

Introduction: Ransomware Has Become a Service Economy

Ransomware was once the product of lone cybercriminals creating rudimentary malware. It's an industry now. On dark web marketplaces, ransomware-as-a-service (RaaS) kits are offered for sale and rental, complete with user guides, revenue-sharing plans, and customer support. By signing up for a "service," even inexperienced attackers can now initiate enterprise-grade ransomware campaigns.

The most hazardous aspect? These operations thrive by taking advantage of privileged access rather than just malware. The "master keys" that attackers require are cloud keys, VPN login credentials, and domain admin accounts. They enable ransomware to spread laterally, encrypt systems at scale, and disable backups once they are stolen or misused.

The lesson for today's executives is straightforward: **RaaS attacks are fueled by privileged access.**

Examples from the Real World: How Privilege Caused Disaster

- **Colonial Pipeline (2021):** DarkSide ransomware affiliates were able to shut down fuel pipelines that served the U.S. East Coast by using a **compromised VPN account** without multi-factor authentication. \$4.4 million was the ransom. The actual loss was a shattered public trust and a national disruption.
- **Healthcare Providers:** Ransomware outbreaks have affected a number of hospitals in the United States, with **hackers gaining access** through compromised service accounts connected to legacy systems. When patient data and devices were rendered inaccessible, lives were literally in danger.
- **Kaseya Supply Chain Breach (2021):** Ransomware was distributed to downstream clients by RaaS groups using remote management tools with privileged access, impacting thousands of small businesses worldwide.

When **RaaS operators** have your privileged accounts, they don't need **zero-day exploits**, as these incidents demonstrate.

Why Conventional Defenses Are Ineffective

- **Perimeter Security Is Obsolete:** Firewalls are useless once an administrator credential is compromised.
- **Inadequate Password Policies:** Phishing, dark web credential dumps, and credential theft make it simple to get around complicated passwords.
- **IAM Is Insufficient:** Identity platforms verify your identity, but they hardly ever identify abnormal behavior from privileged accounts in real time.
- **Fragmented PAM Practices:** Rather than viewing Privileged Access Management (PAM) as a dynamic security discipline, many organizations view it as a compliance checkbox.

RaaS groups take advantage of these flaws by moving swiftly and stealthily within networks, frequently going unnoticed until encryption is activated.

Making the Danger More Human: The Insider and the Forgotten Administrator

Think about a regional bank's system administrator. His privileged account was not appropriately deprovisioned after he left the company two years ago. A RaaS affiliate paid a few dollars to purchase his old credentials when they appeared in a dark web marketplace. Thousands of customers' transactions were frozen as ransomware quickly spread throughout vital banking systems.

Or consider an overburdened IT staff that, for convenience, uses the same "**master admin**" password. One team member is tricked into responding to a phishing email. An attacker only needs that one mistake to take control of backups, turn off monitoring software, and spread ransomware over the weekend of payday.

These reflect innumerable forensic reports and are not merely theoretical.

An Executive Playbook Protecting Privilege from RaaS

1. Reduce the Area of Attack

- Use Just in Time (JIT) access instead of standing privileges.
- Immediately decommission accounts that are not in use.
- Examine and eliminate credentials that are embedded or hardcoded into scripts.

2. Layered Verification

- Enforce multi-factor authentication (MFA) for all privileged accounts, including machine and service identities.
- Employ phishing-resistant techniques, such as hardware tokens and FIDO2.

3. Behavioral Monitoring

- Use AI/UEBA (User & Entity Behavior Analytics) in PAM tool deployment.
- Look for irregularities, such as an administrator logging in from a foreign IP at two in the morning.

4. Divide and contain

- Keep production and admin networks apart.
- Use least privilege and microsegmentation to restrict lateral movement.

5. Recovery and Resilience

- Prevent administrator manipulation of backups (immutable backups).
- Test recovery drills for ransomware every three months.

6. Board Level Oversight

- CISOs are required to report on the risk of privileged accounts to boards on a quarterly basis.
- Consider PAM more than just a compliance tool; treat it as a business continuity control.

Prospects for the Future: Advantage in a RaaS Environment

The RaaS marketplace will continue to advance, using AI to automate privilege escalation and lateral movement. In order to defend privileged access, it will be necessary to implement risk-based access decisions, continuous authentication, and automated anomaly remediation.

Executives need to understand that ransomware cannot be "purchased" with ransom payments or cyber insurance. Proactive privilege governance, which shuts the door before attackers even enter, is the only long-term defense.

Conclusion: Privilege Is the Battlefield

Cybercrime has become a commodity due to **Ransomware as a Service**. Attackers only need your admin passwords; they don't need to outsmart your firewalls. Your privileged accounts are often the weakest link in the chain.

The directive is unambiguous for CEOs, CISOs, and CTOs

- Make PAM a primary resilience strategy investment.
- Require instant access to information about the use of privileged accounts.
- Consider each privileged credential as a possible weapon that ransomware organizations could use.

Privilege is power in the RaaS era. If you don't protect it, attackers will use it as a weapon against your company.

About the Author

Sandeep Dommari is a Senior Cybersecurity Architect and IAM Strategist with over 18 years of experience designing secure access frameworks across Fortune 100 enterprises. His work focuses on application security, adaptive identity, and building secure-by-design architectures for critical industries.

Sandeep can be reached online at sandeep.dommari@ieee.org





Reachability and Exploitability

How to cut through the noise in modern AppSec

By Julia Lorenz, Solutions Manager at Xygeni

Software today is built at a speed and scale we've never seen before. Teams release updates weekly, sometimes daily, and they rarely start from scratch. Instead, modern applications are assembled like an intricate puzzle, part custom-built code, part open-source components, and part third-party packages pulled from public repositories.

This approach powers innovation, but it comes with a hidden cost: you're not just shipping your code; you're shipping the security risks of everyone else's code too. The vulnerabilities of that open-source library you grabbed last year? They're your vulnerabilities now.

To manage this, most organizations rely on [Software Composition Analysis \(SCA\) tools](#). These tools scan your codebase, identify all your dependencies, and compare them against public vulnerability databases such as the National Vulnerability Database (NVD). In theory, that should give you a clear picture of where you're exposed.

But in reality, SCA often produces the same result: a massive, intimidating list of vulnerabilities, with no indication of which ones can be exploited in your environment. The list might be hundreds, even thousands, of entries long. Developers get alert after alert, patch after patch, until they're buried in noise. And somewhere in that pile, there might be one or two vulnerabilities that could truly compromise your system, but they're lost in the flood.

This is the central problem in modern vulnerability management: not every vulnerability is created equal, and most tools don't help you tell the difference.

Why Does Traditional Vulnerability Scanning Fall Short?

Imagine you're running a large e-commerce platform. You push out code every week, your dependency tree is hundreds of packages deep, and your SCA tool just flagged 1,200 vulnerabilities.

Do you fix them all? You can't; it would take weeks of developer time, and in the meantime, your roadmap would grind to a halt. Do you fix only the "critical" ones? That sounds reasonable, until you realize that a low-severity vulnerability in your payment processing code could be far more dangerous than a critical vulnerability in a library you don't even call.

This is where traditional approaches break down: they use severity as the main sorting mechanism. But severity alone is a blunt instrument. It tells you how bad a vulnerability could be in the abstract, not whether it can harm your application as it's deployed today.

The missing piece is context. Without it, teams waste time patching code paths that can never be exploited, while real threats linger in production. This leads to alert fatigue, developer frustration, and a backlog of unaddressed issues.

The Missing Question: Can It Even Run?

This is where **reachability analysis** changes the game.

Instead of simply telling you that a vulnerability exists somewhere in your code or one of your dependencies, SCA reachability analysis asks a much more practical question:

"Can this vulnerable code be reached during my application's runtime execution?"

If the answer is "no," the vulnerability is still worth tracking, after all, future code changes might introduce a path to it, but it doesn't demand immediate attention. If the answer is "yes," then you know it's exploitable in your current setup, and it moves to the top of your list.

By adding this one layer of insight, you transform vulnerability management from a reactionary "patch everything" scramble into a focused, strategic process.

How does SCA Reachability Work in Practice?

Reachability analysis comes in different flavors, each offering a different balance of precision and breadth. Understanding these helps you see why it's such a powerful tool for security teams and developers alike.

The most precise approach is code-level reachability. This method inspects your application's call graph, essentially a map of how functions call one another, and checks whether the vulnerable function is ever invoked, either directly or indirectly. If the function doesn't appear in any execution path, the vulnerability is marked as not reachable.

A classic example comes from a well-known jQuery vulnerability, CVE-2014-6071, which allows remote attackers to conduct cross-site scripting (XSS) attacks via vectors related to use of the `text()` method when used together with `after()`. If your code never uses that exact combination, then the vulnerability cannot be triggered in your environment. Your SCA tool might still flag it, but an SCA reachability analysis would let you confidently deprioritize it.

The broader, less precise approach is **dependency-level reachability**. Here, the question is not "is the vulnerable function called?" but rather "is the dependency itself used at all?" If the answer is yes, it's flagged as potentially reachable. This method doesn't pinpoint exact execution paths, but it can still help you understand whether a vulnerable component plays any role in your application.

Finally, there's the distinction between **always reachable** and **not reachable** vulnerabilities. Some code runs automatically at startup, on every request, or during critical initialization processes. If that code contains a vulnerability, it's inherently high priority because it's executed every time the application runs. By contrast, if a vulnerable function exists in a code path that is never called, it poses no immediate risk, though it should still be monitored for future changes.

Why Exploitability Completes the Picture

Reachability tells you whether the vulnerable code can be executed in your application. Exploitability takes it one step further: is there a practical, real-world way for an attacker to take advantage of it?

This matters because not all reachable vulnerabilities are equally dangerous. A reachable vulnerability that has a known public exploit, especially one that's being actively used in attacks, is an urgent, drop-everything issue. A reachable vulnerability with no known exploit is still important, but its priority might be lower in a high-volume environment where triage is essential.

Severity ratings alone don't reflect this reality. A "low severity" reachable vulnerability in a login endpoint with a known exploit can cause far more damage than a "critical" vulnerability in a rarely used admin feature that's not even exposed to the internet.

By combining reachability with exploitability, you start to see vulnerabilities not as abstract risks but as concrete threats, ranked by their actual potential to cause harm in your specific context.

Traditional Prioritization (Severity-Based)

Severity	Vulnerability	Priority
Critical	Vulnerability in a rarely used feature	High
Low	Vulnerability in a core function	Low

Reachability-Driven Prioritization

Reachability	Vulnerability	Priority
Reachable	Vulnerability in a core function	High
Not Reachable	Vulnerability in a rarely used feature	Low

From Theory to Day-to-Day AppSec

For development and security teams, the shift from severity-based prioritization to reachability-and-exploitability-based prioritization is transformative.

It means fewer false positives cluttering up your backlog. It means less developer time spent chasing vulnerabilities that can't hurt you. And it means faster responses to the vulnerabilities that can.

Organizations that make this shift often report striking improvements. False positives can drop by 70%, remediation times can improve by a third, and developer engagement with security can increase dramatically because fixes feel purposeful instead of arbitrary.

One real-world example comes from Fintonic, a financial services platform. Before adopting reachability-driven prioritization, their security team was buried under thousands of vulnerability alerts, most of them irrelevant. After integrating SCA reachability analysis into their workflow, they reduced false positives by 70% and cut prioritization time by 90%. For them, the change wasn't just about efficiency; it was about regaining control over their security process.

Why This Matters for DevSecOps

In a DevSecOps environment, security can't be a bottleneck. Code moves from commit to production in hours, sometimes minutes. That pace leaves no room for multi-week patch cycles every time a scan runs.

SCA Reachability and exploitability analysis fit naturally into this model. Integrated into the CI/CD pipeline, they can flag exploitable vulnerabilities as soon as they appear, often before the code is even deployed. This keeps security work proactive instead of reactive.

They're also adaptive. As your code changes, as you add features, refactor modules, or swap dependencies, previously unreachable vulnerabilities may become reachable. Continuous reachability analysis means those changes are caught in real time, not during an annual review.

Perhaps most importantly, this approach aligns security work with business priorities. Not every vulnerability affects core services or sensitive data. By ranking vulnerabilities based on both technical exploitability and business impact, teams can focus on protecting what matters most first.

Moving Beyond the Numbers

The security industry has long leaned on numbers, CVSS scores, severity ratings, and vulnerability counts to guide decision-making. But these numbers, while useful, can be misleading in isolation.

A high CVSS score might indicate a dangerous vulnerability *in theory*, but if that code isn't reachable in your application, it's not an urgent fix. Conversely, a low-severity vulnerability might be the easiest path for an attacker to compromise your most critical service.

SCA Reachability and exploitability bring a layer of reality to these numbers. They tell you not just what could go wrong in some hypothetical worst case, but what can go wrong in your application today.

As a Takeaway...

Traditional SCA tools tell you what vulnerabilities exist in your code and dependencies. SCA Reachability analysis tells you whether those vulnerabilities can be executed. Exploitability analysis tells you whether an attacker can realistically take advantage of them. Together, these approaches cut through the noise, focus developer effort where it counts, and dramatically improve both security posture and team morale.

Software moves faster than ever, and the attack surface is constantly shifting; it's no longer enough to know what's broken. You have to know whether it's reachable, whether it's exploitable, and whether it matters right now. That's the difference between chasing vulnerabilities endlessly and securing your applications!


Deep dive [here](#)

Learn more about [reachability-focused security](#)

About the Author

Julia Lorenz is a Solutions Manager at Xygeni. She is a DevSecOps advocate and Software Supply Chain Security specialist. With nearly three years at the company, she has led initiatives to secure clients' software supply chains, embedding security into continuous development cycles and aligning DevOps practices with robust risk mitigation. Her background combines technical expertise with a strong commitment to improving security culture across development teams. A graduate of Technische Universität Berlin, Julia is passionate about raising awareness of software supply chain risks and championing practical, developer-friendly security solutions. Julia can be reached online at julia.lorenz@xygeni.com or via LinkedIn <https://www.linkedin.com/in/julia-e-lorenz/> and at our company website <https://xygeni.io/>





Responsible AI in the Cloud: Building A Framework for Trust

A Blueprint for Developing AI Systems Rooted in Responsibility and Trust

By Marina Bregkou, Principal Research Analyst and Associate Vice President, Cloud Security Alliance

Artificial Intelligence (AI) is now at the heart of cloud computing. From streamlining operations to enabling intelligent decision-making at scale, AI is revolutionizing what's possible in the enterprise. But as AI advances at excessive speed, especially in cloud environments, it has outpaced the rules, safeguards, and ethics we need to use it responsibly.

In the absence of foundational safeguards, AI systems can introduce significant risks: model manipulation, hallucinations, data leakage, adversarial inputs, bias in decision-making, and lack of explainability. These aren't just technical glitches; they're systemic problems that ripple through shared cloud platforms, supply chains, and entire business ecosystems.

As AI scales, our commitment to using it responsibly must scale with it. The foundation of responsible AI in the cloud is shared trust, built through practical frameworks, collaboration across stakeholders, and clear, measurable standards. A safer future for AI is possible, but only if the industry comes together to create and adopt initiatives that make AI secure, ethical, and trustworthy, starting with where it lives: in the cloud.

Laying the Groundwork for Responsible AI

Responsible AI encompasses more than just privacy or fairness: it calls for a full set of ethical, operational, and technical controls that span the entire AI lifecycle. In the cloud, these principles need to align with the shared responsibility model and fit seamlessly into DevSecOps, data governance, and compliance processes.

To move from principles to real-world practice, the industry should focus on four key areas:

- Security and resilience: Protecting AI models and training data from tampering, theft, and misuse
- Transparency and explainability: Making sure cloud-hosted models can be audited, understood, and trusted
- Bias mitigation: Detecting and reducing unfair or discriminatory outcomes
- Governance, accountability, and compliance: Defining clear roles, responsibilities, and ownership across cloud providers and customers.

These are not just ideals. They are now essential building blocks for enterprise trust, regulatory alignment, and AI safety at scale.

Driving Progress Through Global Initiatives

As the leading organization setting standards, certifications, and best practices for secure cloud computing, CSA has played a pivotal role in advancing industry-wide efforts to make AI safe and trustworthy. Through a series of global initiatives, CSA is equipping enterprises, cloud providers, and professionals with the tools and guidance they need to lead confidently in the AI era.

One of the cornerstone tools in this effort is the newly published [AI Controls Matrix \(AICM\)](#), a comprehensive set of security and governance controls mapped specifically to AI use in cloud environments. The AICM helps organizations assess, implement, and validate responsible AI practices against recognized standards, bridging the gap between high-level principles and day-to-day operations.

The [AI Safety Initiative](#), launched in partnership with industry leaders like Amazon, Anthropic, Google, Microsoft, and OpenAI, along with experts from the Cybersecurity & Infrastructure Security Agency (CISA), is a strategic program designed to help organizations deploy AI safely, ethically, and in compliance with emerging standards. Paired with the [AI Safety Ambassador Program](#), it empowers organizations of all sizes to implement AI solutions that are not just innovative, but also responsible and compliant.

The rise of any new technology creates a clear need for a workforce skilled in deploying it safely and securely. The [Trusted AI Safety Knowledge & Certification Program](#) meets that need by providing security, compliance, and risk professionals with the expertise to assess and implement trustworthy AI systems in cloud-first environments.

Organizations that want to show their commitment to responsible AI, in the cloud and beyond, can sign the [AI Trustworthy Pledge](#), the first global agreement calling on cloud and AI stakeholders to uphold principles of transparency, fairness, accountability, and security. The pledge also paves the way for [STAR for AI](#), an upcoming expansion of CSA's [Security, Trust, Assurance, and Risk \(STAR\) Registry](#) to include AI services, and which is expected to launch later this year.

The Role of Cloud Stakeholders

Whether you are a cloud provider, a CISO, an AI developer, or a regulator, you play a critical role in making AI safe, ethical, and trustworthy. Cloud providers need to build responsible AI directly into their platforms, offering secure model deployment, robust data governance, and built-in compliance tools. Cloud customers, in turn, should assess vendors based on AI transparency, fairness, and explainability, making sure SLAs include clear ethical and security commitments.

Meanwhile, security leaders must break down silos and work closely with legal, risk, and development teams to craft AI assurance policies that are proactive, testable, and ready to evolve with emerging threats.

Just like traditional cloud security, responsible AI needs to be built into every layer of the stack, and embedded at every stage of the development pipeline.

From Awareness to Action

AI regulation is gaining momentum. With the EU AI Act, U.S. executive orders, and a growing patchwork of global frameworks, enterprises can no longer treat responsible AI as a side project.


The good news? Across the industry, organizations are coming together to deliver the tools, frameworks, and community support needed to make this transformation possible. But real progress takes action. Organizations must commit, not just in principle, but also in practice, to building AI systems that are both innovative and trustworthy. That commitment begins with a strong foundation of cloud security, shared accountability, and a clear goal: ensuring AI serves people, not just profit.

AI in the cloud is transforming the digital world. But without a strong foundation of responsibility and trust, its promise can just as easily become a risk. By adopting collaborative, standards-driven approaches and supporting initiatives that make AI safe and accountable, the industry can move beyond reactive compliance, and step into true, proactive leadership.

Responsible AI in the cloud is not just possible, it is essential, and the time to build it is now.

About the Author

Marina Bregkou is the Principal Research Analyst of the Cloud Security Alliance. She holds a BSc of Informatics and Telecommunications from the Athens National University, and a MSc of Information Systems from the Athens University of Economics and Business. Since starting at CSA, she has been involved in various FP7 and H2020 projects dealing with cybersecurity, cloud computing, information security, cloud interoperability, security standards, service level agreements, etc. Beginning in 2019, Marina has been a lead and contributing author to at least 15 CSA artifacts focusing on containers, microservices architecture, serverless computing and architecture, the management of keys in the cloud and artificial intelligence. She also manages and directs the relevant CSA working groups. Her professional and scientific interests, among others, include privacy in the cloud and education of senior adults in the secure use of the internet, its tools and applications. Marina can be reached online at mbregkou@cloudsecurityalliance.org and at <https://cloudsecurityalliance.org/>.





SASE Architecture Deployment Across the Globe

By Vishal Gudhka, Senior Network Architect, Versa Networks

Enterprises are rapidly transforming their networking by deploying databases, web, storage, and AI applications in both on-premises data centers and the cloud. The main reason for this transformation is to meet the needs of a hybrid workforce that requires secure and seamless access to applications from anywhere in the world. In doing so, they aim to avoid network latency and protect against various cyberattacks. The Enterprises deployed traditional solutions that provide disparate security and network functionality. Below are the challenges Enterprises face.

Challenges of Siloed and Traditional Solutions

1. **Disintegrated Architecture** – Using different vendor solutions for security and networking separately. These separate solutions lead to inefficient routing paths and expose security vulnerabilities.

2. **Erratic Application Access** – Users experience latency and increased internet risks when accessing applications remotely hosted in the cloud or a data center.
3. **Inconsistent Policies** – The inconsistent configuration policies across different sites lead to routing and security performance issues. Many inbound or outbound malicious activities aren't blocked, and a latency is experienced while accessing corporate sensitive applications.
4. **Access to all corporate data** – Traditional solutions, such as VPN, are perimeter-based and grant full access to the enterprise's data once a user is logged in.
5. **Insufficient visibility resulting in operational inefficiencies** – There was no single source of truth for user behavior analytics, device posture assessment, or security event correlation.
6. **Operational Inefficiencies** – When incidents occurred, the operations team had to manually sift through disjointed logs from multiple systems, resulting in a prolonged mean time to resolution (MTTR).

Secure Access Service Edge (SASE) Solution

A well-thought-out and designed SASE solution addresses all the enterprise's demands.

SASE Gateways

- **Cloud-Native Architecture** – Cloud-native architecture is easily scalable, allowing for horizontal or vertical scaling based on demand. This approach ensures a design with high availability and fault tolerance across all the SASE gateways. This architecture also provides the following benefits.
- **Single Software Architecture** – The unified software solution provides integrated networking, including Software-Defined Wide Area Network (SDWAN), and security service edge (SSE) features, such as Secure Web Gateway (SWG), Cloud Access Security Broker (CASB), Zero Touch Network Access (ZTNA), and Firewall-as-a-Service (FWaaS). The traffic passing through the software architecture is analyzed only once, which improves performance, reduces processing overhead, and lowers the risk of security vulnerabilities.
- **Zero Trust Principle** – Users are granted the least privileged access, allowing them to view only the data assigned to their role. Throughout the network, users and their devices are continuously authenticated and assessed before access is permitted.
- **Multitenancy** – Enterprises run multiple tenants to keep important and sensitive data traffic separate from public traffic, such as hosting a public tenant for WIFI for guests and non-employees. Additionally, unique traffic steering and security policies are tailored to specific business needs. A significant benefit for enterprises is the ability to quickly create an extra lab tenant to test new configuration policies and software versions before deploying them in production.
- **Consistent Policy Enforcement** – The IT administrator can create one or multiple policies and implement them across all SASE gateways. These policies can now include networking services configurations, such as application-aware or SLA-based routing, DHCP, DNS, QoS, VLANs, VRRP, and security services like URL filtering, IP filtering, file filtering, DNS filtering, malware scanning, anti-virus, Data Loss Prevention (DLP), web and cloud access control, and more. A

policy with specific configurations can be assigned to selected SASE gateways to meet business requirements alongside a policy with common configurations.

Centralized Orchestration and Data Lake

The cloud-native orchestration and Data Lake platforms are essential in SASE architecture.

Data Lake Platform

All SASE gateways send the statistics, security logs, and alarms to a centralized data lake platform. The data lake platform is integrated with User and Entity Behavior Analytics (UEBA) that correlates all logs into actionable threat intelligence. Troubleshooting becomes easier as the operations team can quickly review the logs or reports to determine the root cause and apply the workaround or direct fix. Machine Learning (ML) is also integrated to further enhance the capabilities of MLOps. The platform now baseline the user and device behavior. Any deviation from this behavior would notify the operations teams to take corrective steps promptly. This further reduces the MTTR and MTTR.

Orchestrator Platform

The orchestrator platform simplifies the work for operations and design teams. The teams can easily

1. Create policies through workflow and apply them to all SASE gateways at once.
2. View any gateway that is out of sync from the configuration perspective.
3. Upgrade, downgrade, and rollback configurations to minimize service disruptions.
4. Navigate to the analytics tab to analyze logs and generate reports for the business teams.

Reduction in Cost

SASE significantly lowers CapEx and OpEx because enterprises no longer need to deploy multiple solutions. SASE provides comprehensive SDWAN and SSE features on a single platform. SDWAN replaces MPLS circuits with broadband, reducing WAN costs substantially. It also offers networking services like DHCP, QoS, Multitenancy, routing, and switching. SSE features include SWG, CASB, FWaaS, DLP, and ZTNA. Operational expenses are also decreased since enterprises can utilize an orchestration platform for automation and policy management, as well as a Data lake platform to decrease MTTR and MTTR. The entire solution is cloud-native, resulting in significant cost reductions in power costs, and a smaller number of technicians is required during an outage situation.

SASE Client Registration and Gateway Connection

The SASE client is an agent used to connect to the SASE gateway. The first phase is for the client to register with the gateway. During registration, the user authenticates through Single Sign-On or Multi-

Factor Authentication (MFA). The device posture check is also performed during this phase, where corporate guidelines are evaluated against users' endpoints to ensure OS versions, anti-virus software, disk encryption, and other security measures are running on the endpoint. On successful authentication, the client downloads the user-specific and device-specific configuration profile. The configurations include,

1. **Routing Policies** – Specific routes to the applications in the cloud or data center are downloaded to ensure the shortest path to applications. Sometimes, direct routing between the SD-WAN devices is enabled to reach applications hosted at specific sites directly.
2. **DNS Settings** – The split DNS feature improves performance, lowers latency, and boosts security. Resolve all internal domain names using the internal enterprise DNS server. This allows quick access to internal applications and prevents internal domain names from being exposed to external DNS servers. To offload the internal DNS infrastructure, resolve public domain names through trusted cloud public DNS servers.
3. **Device Location** – The client's location is crucial so that the correct information is downloaded and users are connected to the closest gateway, thereby optimizing performance.

The next phase is gateway connectivity, where the client sends probes to gather real-time metrics, such as gateway proximity and load. This helps the client establish a secure connection with a well-functioning SASE gateway. For example, multiple gateways are deployed across the East Coast, Central, Midwest, and West Coast of the US. Users in Chicago should connect to the Central gateway, as it is likely the closest to users. If the gateway in Central is not healthy, the device may connect to the next closest gateway, which is likely an East Coast gateway. Device posture checks are periodically performed, even after the device is connected to the gateway, to uphold zero-trust policies and ensure business compliance.

TLS Certificate Management:

A well-known public key infrastructure is used to secure control and data traffic between the SASE client and gateways. Enterprises can choose to use a single certificate for all SASE gateways or a separate certificate for each gateway. The recommendation is to deploy the latter approach as it is considered more secure. Enterprises can leverage the orchestration platform to automate the lifecycle of certificate management, making operations more straightforward.

Data Compliance

Data laws vary across the globe. Many countries have stringent laws that make it invalid to transfer and process data beyond their borders. Here, SASE infrastructure becomes a significant factor in supporting these countries' laws. Traffic is locally routed to ensure compliance and not compromise network performance.

Conclusion

As mentioned above, A meticulous SASE design on a global scale offers a great user experience. The network's performance will be optimal and secure from outside attacks. The design also simplifies the day-to-day workload of operations and design teams, often empowering them to establish a strong network that complies with Zero Trust principles.

About the Author

Vishal Gudhka, Senior Network Architect at Versa Networks, has expertise in SDWAN and SASE. Vishal has over 14 years of experience in service provider, data center, and cloud networking technologies. At Versa Networks, Vishal has led transformative projects, working with customers to modernize their networks. He has architected a scalable SASE infrastructure with enhanced security, improved network performance, and significantly reduced deployment and operational costs. In recognition of his achievements, he also earned a specialist-level certification in SD-WAN. Previously, Vishal served as a Senior Resident Engineer at Juniper Networks, where he designed and deployed complex, high-performance network infrastructures. Over six years, he obtained three expert-level certifications: JNCIE-SP, JNCIE-ENT, and JNCIE-DC. He has earned a master's degree in electrical engineering from The George Washington University. Combining his master's academic depth with practical expertise in developing secure, scalable network solutions. Vishal Gudhka can be reached online at gudhka.vishal@gmail.com or <https://www.linkedin.com/in/gudhkavishal>.





Security Chaos Engineering for CISOs - The Strategic Edge Against Modern Threats

By Andres Andreu, CISO and COO, Constella Intelligence

In an age where cyber attackers have become more intelligent, agile, persistent, sophisticated, and empowered by Artificial Intelligence (AI), defenders must go beyond traditional detection and prevention. The traditional models of protective security are fast becoming diminished in their effectiveness and power. In the face of pursuing a proactive model one approach has emerged, security chaos engineering. It offers a proactive strategy that doesn't just lead to hardened systems but can also actively disrupt and deceive attackers during their nefarious operations.

By intentionally injecting controlled failures or disinformation into production-like environments, defenders can observe attacker behavior, test the resilience of security controls, and frustrate adversarial campaigns in real time.

What is Security Chaos Engineering?

Security chaos engineering is the disciplined practice of simulating security failures and adversarial conditions in running production environments to uncover vulnerabilities and test resilience before

adversaries can. Its value lies in the fact that it is truly the closest thing to a real incident. Table Top Exercises (TTXs) and penetration tests always have constraints and/or rules of engagement which distance them from real world attacker scenarios where there are no constraints. Security chaos engineering extends the principles of chaos engineering, popularized by Netflix (<https://netflixtechblog.com/chaos-engineering-upgraded-878d341f15fa>) to the security domain.

Instead of waiting for real attacks to reveal flaws, defenders can use automation to introduce “security chaos experiments” (e.g. shutting down servers from active pools, disabling detection rules, injecting fake credentials, modifying DNS behavior) to understand how systems and teams respond under pressure.

Here are some security chaos engineering techniques to consider as this becomes part of a proactive cybersecurity strategy:

Temporal Deception - Manipulating Time to Confuse Adversaries

Temporal deception involves distorting how adversaries perceive time in a system (e.g. injecting false timestamps, delaying responses, or introducing inconsistent event sequences). By disrupting an attacker’s perception of time, defenders can introduce doubt and delay operations.

Example: Temporal Deception through Delayed Credential Validation in Deception Environments

Consider a deception-rich enterprise network, temporal deception can be implemented by intentionally delaying credential validation responses on honeypot systems. For instance, when an attacker attempts to use harvested credentials to authenticate against a decoy Active Directory (AD) service or an exposed RDP server designed as a trap, the system introduces variable delays in login response times, irrespective of the result (e.g. success, failure). These delays mimic either overloaded systems or network congestion, disrupting an attacker’s internal timing model of the environment. This is particularly effective when attackers use automated tooling that depends on timing signals (e.g. Kerberos brute-forcing or timing-based account validation). It can also randomly slow down automated processes that an attacker hopes completes within some time frame.

By altering expected response intervals, defenders can inject doubt about the reliability of activities such as reconnaissance and credential validity. Furthermore, the delayed responses provide defenders with crucial dwell time for detection and the tracking of lateral movement. This subtle manipulation of time not only frustrates attackers but also forces them to second-guess whether their tools are functioning correctly or if they’ve stumbled into a monitored and/or deceptive environment.

Honey Timing and Time-Based Traps

Time-bound honeypots such as fake cron jobs, scheduled updates, or bogus backup routines can serve as deceptive traps. Interaction with these elements reveals unauthorized probing or access attempts.

The very existence of these traps implies that any entity interacting with them (excluding the creators of course) needs to be treated as hostile and investigated.

Example: Deceptive Backup Scripts as Time-Based Traps in Cloud Environments

Defenders can deploy a bogus scheduled backup script named “nightly-db-backup.sh” on a decoy cloud instance. The script can be set to appear as if it ran daily at 04:00 using a convincingly sounding cron job (e.g. /etc/cron.d/backup_job). The script can contain clear-text references to fake database credentials, S3 storage paths, and mock sensitive data exports. This can be used as a timing-based honeypot, existing to attract unauthorized access attempts during off-hours when legitimate activity is minimal.

Any attempt to execute this script triggers hidden canary tokens that act as an alerting system. This can trigger things like an HTTP request where the receiving entity (e.g. web server processing the request) has been configured to log and alert on any relevant interaction. This can of course capture timestamps showing interactions with the script outside of the bogus scheduled execution window. The defenders can then not only detect the unauthorized access but also track subsequent movements due to some of the meta-data captured.

This approach demonstrates how time-based decoy elements, especially those aligned with off-hour routines, can effectively expose stealthy adversaries who are mimicking typical system administrator behavior.

Randomized Friction

Randomized friction aims at increasing an attacker's work factor, in turn increasing the operational cost for the adversary. Introducing unpredictability in system responses (e.g. intermittent latency, randomized errors, inconsistent firewall behavior) forces attackers to adapt continually, degrading their efficiency and increasing the likelihood of detection.

Example: Randomized Edge Behavior in Cloud Perimeter Defense

Imagine a blue/red team exercise within a large cloud-native enterprise. The security team deploys randomized friction techniques on a network segment believed to be under passive recon by red team actors. The strategy can include intermittent firewall rule randomization. Some of these rules make it so that attempts to reach specific HTTP based resources are met with occasional timeouts, 403 errors, misdirected HTTP redirects, or to simply give an actual response.

When the red team conducts external reconnaissance and tries to enumerate target resources, they experience inconsistent results. One of their obvious objectives is to remain undetected. Some ports appeared filtered one moment and opened the next. API responses switch between errors, basic authentication challenges, or other missing element challenges (e.g. HTTP request header missing). This forces red team actors to waste time revalidating findings, rewriting tooling, and second-guessing whether their scans were flawed or if detection had occurred.

Crucially, during this period, defenders are capturing every probe and fingerprint attempt. The friction-induced inefficiencies increase attack dwell time and volume of telemetry, making detection and attribution easier. Eventually, frustrated by the lack of consistent telemetry, the red team escalates their approach. This kills their attempts at stealthiness and triggers active detection systems.

This experiment successfully degrades attacker efficiency, increases their operational cost, and expands the defenders' opportunity window for early detection and response, all without disrupting legitimate internal operations. While it does take effort on the defending side to set all of this up, the outcome would be well worth it.

Ambiguity Engineering

Ambiguity engineering aims to obscure the adversary's mental model. It is the deliberate obfuscation of system state, architecture, and behavior. When attackers cannot build accurate models of the target environments, their actions become riskier and more error-prone. Tactics include using ephemeral resources, shifting IP addresses, inconsistent responses, and mimicking failure states.

Example: Ephemeral Infrastructure and Shifting Network States in Zero Trust Architectures

A SaaS provider operating in a zero trust environment can implement ambiguity engineering as part of its cloud perimeter defense strategy. In this setup, let's consider a containerized ecosystem that leverages Kubernetes-based orchestration. This platform can utilize elements such as ephemeral IPs and DNS mappings, rotating them at certain intervals. These container hosted backend services would be accessible only via authenticated service mesh gateways, but appear (to external entities) to intermittently exist, fail, or timeout, depending on timing and access credentials.

Consider the external entity experience against a target such as this. These attackers would be looking for initial access followed by lateral movement and service enumeration inside this target environment. What they would encounter are API endpoints that resolve one moment and vanish the next. Port scans would deliver inconsistent results across multiple iterations. Even successful service calls can return varying error codes depending on timing and the identity of the caller. When this entity tries to correlate observed system behaviors into a coherent attack path, they would continually hit dead ends.

This environment was not broken; it was intentionally engineered for ambiguity. The ephemeral nature of resources, combined with intentional mimicry of common failure states, would prevent attackers from forming a reliable mental model of system behavior. Frustrated and misled, their attack chain will slow, errors will increase, and their risk of their detection will rise. Meanwhile, defenders can capture behavioral fingerprints from the failed attempts and gather critical telemetry for informed future threat hunting and active protection.

Disinformation Campaigns and False Flag Operations

Just as nation-states use disinformation to mislead public opinion, defenders can plant false narratives within ecosystems. Examples include fake internal threat intel feeds, decoy sensitive documents, or impersonated attacker TTPs designed to confuse attribution.

False flag operations are where an environment mimics behaviors of known APTs. The goal is to get one attack group to think another group is at play within a given target environment. This can redirect adversaries' assumptions and deceive real actors at an operational stage.

Example: False Flag TTP Implantation to Disrupt Attribution

Consider a long-term red vs. blue engagement inside a critical infrastructure simulation network. The blue team defenders implement a false flag operation by deliberately injecting decoy threat actor behavior into their environment. This can include elements such as:

- Simulated PowerShell command sequences that mimic APT29 (<https://attack.mitre.org/groups/G0016/>) based on known MITRE ATT&CK chains.
- Fake threat intel logs placed in internal ticketing systems referring to OilRig or APT34 (<https://attack.mitre.org/groups/G0049/>) activity.
- Decoy documents labeled as "internal SOC escalation notes" with embedded references to Cobalt Strike Beacon callbacks allegedly originating from Eastern European IPs.

All of these artifacts can be placed in decoy systems, honeypots, and threat emulation zones designed to be probed or breached. The red team, tasked with emulating an external APT, stumble upon these elements during lateral movement and begin adjusting their operations based on the perceived threat context. They will incorrectly assume that a separate advanced threat actor is and/or was already in the environment.

This seeded disinformation can slow the red team's operations, divert their recon priorities, and lead them to take defensive measures that burn time and resources (e.g. avoiding fake IOC indicators and misattributed persistence mechanisms). On the defense side, telemetry confirmed which indicators were accessed and how attackers reacted to the disinformation. This can become very predictive regarding what a real attack group would do. Ultimately, the defenders can control the narrative within an engagement of this sort by manipulating perception.

From Fragility to Adversary Friction

Security chaos engineering has matured from a resilience validation tool to a method of influencing and disrupting adversary operations. By incorporating techniques such as temporal deception, ambiguity engineering, and the use of disinformation, defenders can force attackers into a reactive posture. Moreover, defenders can delay offensive objectives targeted at them and increase their attackers' cost

of operations. This strategic use of chaos allows defenders not just to protect an ecosystem but to shape adversary behavior itself.

About the Author

Andres Andreu serves as both the Chief Operating Officer (COO) and Chief Information Security Officer (CISO) at Constella Intelligence. He is a 4X CISO and distinguished cybersecurity leader with credentials including CISSP, ISSAP, and Boardroom Certified Qualified Technology Expert (QTE). His diverse career spans federal law enforcement, where he earned three U.S. Department of Justice awards for contributions to lawful intercept technology, corporate leadership at Hearst, Ogilvy & Mather and 2U, Inc./edX, and entrepreneurial success as a founding executive at Bayshore Networks (acquired by Opswat in 2021). Recognized as a Top 100 CISO (C100) and a Top 50 Information Security Professional, he balances offensive and defensive cybersecurity strategies with a leadership philosophy that aligns executive and employee objectives. An acclaimed author of *The CISO Playbook* and *Professional Pen Testing Web Applications*, he also holds patents in cybersecurity innovations and advises at Forgepoint Capital's Cybersecurity Advisory Council.



Andres can be reached online at [LinkedIn](#) and at our company website <https://constella.ai/>



Generative AI Security: Protecting AI Workloads with AWS Cloud

Building Multi-Layered Security for AI workloads

By Jatinder Singh, Senior Technical Account Manager, Amazon Web Services

The rapid adoption of artificial intelligence (AI) and machine learning (ML) in cloud environments has introduced significant new security challenges. Generative AI, with its ability to create, manipulate, and analyze vast amounts of data, demands robust protection mechanisms.

Based on my experience working with global banking clients since the initial release of Amazon Bedrock, I've seen firsthand how AWS addresses these challenges through an integrated security framework spanning multiple services: Amazon Bedrock, SageMaker, and Amazon Q. While many organizations initially focus on model capabilities, the real challenge lies in securing these AI workloads without

sacrificing performance – a balance I’ve helped numerous financial institutions achieve. This article examines the Defense in Depth approach to securing AI applications, covering data protection, model security, and enterprise integration.

Modern AI workloads face critical security risks including:

- Prompt injection attacks – a threat I’ve seen increasingly targeted at chatbots
- Model output manipulation leading to hallucinations challenging for critical applications
- Training data poisoning attempts
- Model theft through extraction attacks that is a growing concern for proprietary trading models

Drawing from my recent experience helping a G-SIB secure their generative AI infrastructure, organizations follow NIST’s Defense in Depth (DiD) methodology to protect assets across all layers- from infrastructure to model endpoints. What’s particularly effective about this multi-layered strategy is how it combines robust security controls with native safeguards embedded within AWS Generative AI services.

From my work I’ve learned that organizations can confidently advance their AI initiatives while upholding stringent standards for data protection, compliance, and operational security by following a methodical approach tailored to their risk profile.

AWS Defense in Depth Strategy for Security

While Defense in Depth (DiD) isn’t a new concept, its application to AI workloads requires specific considerations. Let me share a real example: When implementing Bedrock for a major investment bank, we discovered that traditional DiD approaches needed significant adaptation to handle the unique challenges of foundation models.

Defense in Depth (DiD) employs multiple layers of security controls to protect data, applications, and infrastructure. What makes this especially crucial for AI workloads is the dynamic nature of model interactions. For instance, one of my clients discovered that traditional data loss prevention (DLP) solutions were insufficient for catching sensitive data in model outputs - we had to implement additional semantic analysis layers.

Each security layer provides distinct but complementary controls. Drawing from my experience here’s a breakdown of how these layers work together in practice:

Security Layer	Primary Function	Key AWS Security Services
Data Protection (Core Layer)	Protect data at rest and in transit	AWS KMS AWS CloudHSM Amazon Macie AWS Secrets Manager

		S3 Object Lock Amazon S3 Encryption
Application Protection	Secure applications and APIs	AWS WAF AWS Shield AWS Certificate Manager AWS Verified Access
Infrastructure Protection	Secure compute and container resources	Amazon Inspector AWS Security Hub EC2 Security Groups AWS Systems Manager
Network & Edge Protection	Network security and connectivity	AWS Network Firewall Amazon VPC AWS PrivateLink AWS Transit Gateway
Threat Detection & Response	Monitor and respond to threats	Amazon GuardDuty Amazon Detective AWS Security Hub AWS CloudTrail Amazon CloudWatch
Identity & Access Management	Manage identities and access	AWS IAM IAM Identity Center Amazon Cognito AWS Directory Service AWS Organizations AWS Control Tower

		IAM Access Analyzer AWS Verified Permissions
Policies & Procedures	Governance and compliance	AWS Audit Manager AWS Organizations (SCPs) AWS Config AWS Trusted Advisor AWS Well-Architected Tool

Securing Generative AI services on AWS

With this Defense in Depth framework in place and all services providing layered security, let's explore how AWS services augment the security across its generative AI services. Based on my implementation experience, three Generative AI services are particularly notable to be discussed from this perspective:

1. Amazon Bedrock

Amazon Bedrock has evolved into AWS's cornerstone for secure foundation model deployment. Its security features address the core challenges regularly encountered in enterprise implementations:

Data Privacy

- No customer data used for model training
- Private customization capabilities within customer VPC
- Secure model fine-tuning with customer-controlled encryption
- Enhanced PII detection capabilities

Access Controls

- Fine-grained IAM permissions for model access
- Role-based controls for model management
- Secure API authentication

Network Security

- VPC Endpoint support
- Private network connectivity
- TLS 1.2 encryption minimum for data transit

Content Safety

- Configurable content filtering

- Custom topic restrictions
- Automated content monitoring

Bedrock integrates with core AWS security services to provide:

Encryption

- AWS KMS integration for data at rest
- TLS 1.2+ encryption in transit
- Customer-managed key support

Network

- VPC endpoint integration
- AWS PrivateLink support
- Granular access controls

Monitoring

- CloudWatch metrics and logging
- CloudTrail API activity tracking
- Detailed security event monitoring

Compliance

- SOC 1, 2, & 3 certifications
- HIPAA eligibility
- PCI DSS compliance
- GDPR controls

2. Amazon SageMaker

Amazon SageMaker's security framework has significantly evolved since the service has been launched and below are some features that prove most effective in production environments:

Model Security

- Secure training environments with isolated compute
- End-to-end artifact encryption
- Advanced model tampering detection
- Continuous model behavior monitoring

Data Protection

- Secure Data Wrangler workflows
- Protected Ground Truth labeling
- Enhanced pipeline security controls

- Automated PII detection in training data

Governance

- Model Cards for documentation and lineage
- Role Manager for access control
- Clarify for bias detection
- Advanced model explainability features

Key security integrations we commonly implement:

Access Management

- SageMaker Role Manager
- Fine-grained IAM controls
- Resource-level permissions
- Cross-account model sharing controls

Monitoring Capabilities

- Model Monitor for drift detection
- Automated endpoint monitoring
- AI-specific threat detection
- Real-time security alerts

Network Controls

- VPC support
- Private endpoints
- Network isolation
- Enhanced traffic monitoring

Governance Features

- Model Cards
- Lineage tracking
- Audit capabilities
- Compliance reporting

These controls have proven particularly valuable in regulated industries where model governance is critical.

3. Amazon Q

Amazon Q being a generative AI-powered assistant provides conversational business intelligence, enabling organizations to have natural language interactions with their data and systems while

maintaining enterprise-grade security. Here's what matters most in production from Amazon Q's implementation perspective:

Access Management

- Role-based controls aligned with organizational structure
- Enterprise authentication integration
- Permission inheritance across teams
- Context-aware access patterns

Data Security

- Zero model training on customer content
- Secure document processing
- Granular data access controls
- Enhanced PII protection mechanisms

Enterprise Integration

- SSO provider support
- Permission boundary enforcement
- Encrypted connections
- Custom security policies

Real-world security controls we typically implement:

Authentication

- Enterprise identity integration
- Multi-factor authentication
- Behavioral authentication patterns
- Session security controls

Authorization

- Document-level permissions
- Role inheritance
- Dynamic access adjustments
- Granular API controls

Data Protection

- End-to-end encryption
- Secure document handling
- Enhanced audit logging
- Access tracking

Compliance

- Continuous monitoring
- Security reporting
- Audit trails
- Industry-specific controls

These security features create a protected environment for AI-powered business operations while maintaining compliance requirements.

Conclusion

Based on my experience implementing these security frameworks across multiple enterprises, it's clear that AWS's defense-in-depth approach provides a robust foundation for secure AI operations. Let me summarize the key takeaways:

AWS's defense-in-depth strategy provides comprehensive protection for AI workloads, demonstrated through key services. Amazon Bedrock has proven robust in handling sensitive data, with strong model access controls and effective protection against prompt attacks. Amazon SageMaker delivers end-to-end model security with robust monitoring capabilities and strong compliance controls, while Amazon Q provides enterprise-grade security with effective data privacy and seamless security integration.

Looking ahead, organizations should prepare for new AI regulations, implement continuous monitoring, stay current with emerging threats, and maintain strong access controls. This approach helps organizations innovate with AI while maintaining data privacy, compliance, and security.

About the Author

Jatinder Singh is the Senior Technical Account Manager at Amazon Web Services (AWS). He has over two decades of experience in enterprise architecture and digital transformation. As a veteran architect who has led enterprise-scale modernization programs for global banking corporations and U.S. Global Systemically Important Banks (G-SIBs), he brings deep insights into securing complex AI workloads. His expertise spans cloud-native architectures, data lake technologies, and AI-driven systems, with significant experience in architecting secure, scalable enterprise systems that achieve measurable operational improvements. Jatinder specializes in implementing Defense in Depth strategies while enabling innovation, making him uniquely qualified to address the complex security challenges facing organizations adopting AI technologies.



Jatinder can be reached online at jatinder.singh.work@gmail.com and on LinkedIn at <https://www.linkedin.com/in/jatinderaws>



Part 1: Maximizing the Value of LLMs Without Compromising Security: The Identity Risk Behind AI Agents

Why AI Agents Pose a Growing Threat to Identity Security

By Amit Zimmerman, Co-Founder and CPO, Oasis Security

Large language models (LLMs) such as ChatGPT, Claude, and Llama are transforming the way organizations conduct, manage and grow their business. Just one year after OpenAI released ChatGPT, OpenAI CEO Sam Altman [shared](#) that over two million developers use the platform, including more than 92% of Fortune 500 companies.

The [model context protocol \(MCP\) published by Anthropic](#) in November 2024 opened the doors for even more innovation through seamless integration. Thousands of MCP servers are published online for developers to utilize, and the MCP repository has already been forked over 4,000 times.

These tools are being connected and integrated into organizations' environments and bring with them novel risks to [Non-Human Identity](#) (NHI) security, access to data, and resource entitlements. Against this backdrop, security professionals with a "should I allow this?" approach must consider adopting a "how do we secure this?" approach.

Security Considerations as AI Becomes Increasingly Pervasive in the Workplace

Leveraging GenAI solutions requires connecting them to your data. Whether it is to analyze trends in usage data, or automate and streamline your user experience, connectivity is key. And the ways to connect models to your environment is growing as well. For example, you can now connect an AI agent to a plethora of services your organization already uses. However, these agents don't just analyze data; they can take active actions and execute automations across systems. This opens up powerful new possibilities, but also significantly increases the potential blast radius of a bad actor or a misconfiguration.

This connectivity is facilitated by creating new identities or leveraging existing ones. In both cases, there is a risk of over-provisioning permissions or assigning roles that are poorly suited for the task, especially if those identities aren't properly monitored or secured.

Some of the most common risks include:

- Over-permissioned access granted to AI agents that may not need broad or persistent privileges
- Unmonitored identities, often created through OAuth flows or personal accounts, that operate without central visibility
- Shadow integrations introduced by individual teams or business units without security oversight
- Poor revocation processes for short-lived or experimental tools that retain access longer than intended

This problem is further compounded by how these integrations are introduced. Developers, eager to explore and implement the latest technologies, may hastily onboard AI tools without fully considering the scope of access being granted. In other cases, non-developers might connect AI services using their own credentials, creating unmanaged backdoors into sensitive environments.

As AI becomes more embedded in everyday workflows, security professionals must ensure that the convenience of integration does not come at the cost of visibility and control.

Treat AI Agents as High-Risk Actors

AI agents are being granted access to internal systems, from CRMs and cloud infrastructure to financial tools and code repositories, and these workloads often operate with elevated privileges—yet there's minimal monitoring of what they're doing with that access. And unlike traditional analytics tools, AI agents can take action: opening tickets, modifying configurations, updating records, or triggering workflows. This level of autonomy makes them powerful but also risky.

For example, within an MCP environment, an AI agent might monitor system usage patterns and, upon detecting performance degradation, proactively open a ticket in your ITSM platform, scale out infrastructure via your cloud provider's API, and update an internal dashboard to reflect the change, all without human intervention. This kind of automation is powerful; but consider the permissions that were granted to facilitate this ability: read production environments logs and metrics; write access to ITSM; modify production deployment configurations.

Now consider that in the rush to implement and utilize, many would simply give free reign to the agent by providing their own credentials, or provisioning access through largely over-permissive identities. It is increasingly difficult to navigate modern IAM requirements. What's more concerning is that these agents often operate in the background, in bulk, and outside of standard security controls. Their activity is under-logged and their behavior remains essentially unmonitored.

Giving AI agents the keys to the kingdom without visibility or guardrails could have severe security consequences. Organizations should treat these identities as high-risk actors, applying least privilege, robust logging, and constant oversight to ensure safety and control.

Coming Next


AI is moving fast, but if identity isn't part of the plan, risk moves faster. Security leaders need to know where AI is running, what it can access, and who—or what—it really is.

Part Two, for future publication, takes it further, diving into how to bake security into AI automation from the start, including best practices for managing machine identities and protecting the credentials that power these tools.

About the Author

Amit Zimerman is the Co-Founder and Chief Product Officer at Oasis Security. He is a seasoned leader with a diverse technical and product background. Before co-founding Oasis, he played pivotal roles at CyberMDX, and Microsoft, bringing a wealth of product and security expertise. Amit also had significant contributions during his seven-year tenure in Israeli Military Intelligence forces as a leader of some of the high-profile cyber projects at the time.

Amit can be reached online at Oasis@icrinc.com and at our company website <https://www.oasis.security/>





The Ghost in the Firewall

How Autonomous AI Will Redefine Cyber Defense

By Bekir Tolga TUTUNCUOGLU, CEO & Researcher, TTNC Teknoloji

For decades, cybersecurity has been a battle of attrition — defenders patching, attackers probing, both sides locked in an endless arms race. In recent years, artificial intelligence has tilted the balance, giving defenders unprecedented speed in detecting anomalies and isolating threats. But beneath the surface, something more radical is brewing: **AI systems that don't just detect attacks — they anticipate, intercept, and neutralize them before they even begin.**

We are on the verge of what could be the most profound transformation in cybersecurity since the birth of the firewall: **autonomous, predictive AI defense agents that act more like living organisms than static software.**

From Reactive Walls to Proactive Predators

Traditional defenses, even AI-enhanced ones, are still largely reactive. They wait for a suspicious event, analyze it, and respond. But the future will belong to **proactive AI guardians** — systems that continuously scout the digital perimeter, model potential adversary strategies, and preemptively shut down avenues of attack before a malicious packet even touches the network core.

Imagine an AI that:

- Maps every dependency and communication channel in real time, identifying new attack surfaces the instant they appear.
- Monitors global cyber threat intelligence feeds, automatically adjusting internal defenses to neutralize tactics used in attacks happening continents away.
- Simulates thousands of potential breach scenarios every hour, using adversarial reasoning to close gaps before they're discovered by human attackers.

This isn't "threat detection" anymore — it's **threat eradication at machine speed**.

The Evolutionary Leap: Self-Learning Cyber Guardians

The biggest leap forward is the shift from static AI models to **self-evolving architectures**. These systems learn from every interaction, refining both their defensive playbook and their understanding of adversarial behavior.

A static AI model is like a locked encyclopedia. A self-evolving AI is like a predator learning new hunting tactics with every encounter. The longer it's deployed, the more cunning it becomes — and the harder it is for attackers to predict or outmaneuver.

Such systems are being trained not just on historical data, but on **synthetic attack simulations** generated by other AIs. This adversarial co-training means the defender AI learns in a high-intensity environment of constant, evolving challenge — the digital equivalent of a black ops training camp.

The Coming Age of Autonomous Incident Response

Today's security teams are plagued by alert fatigue, incident queues, and slow remediation cycles. Autonomous AI promises to collapse that entire process into seconds:

- Detect anomaly.
- Verify threat potential.
- Execute containment strategy.
- Roll out global patches.

All without a single human approval.

This level of autonomy will fundamentally change the role of human analysts — from front-line responders to strategic supervisors, focusing on oversight, compliance, and the development of new detection logic.

A Symbiotic Defense Model

The most effective implementations won't replace humans — they'll form **AI-human hybrid defense ecosystems**. Humans excel at creative reasoning, ethical decision-making, and contextual judgment. AI excels at relentless vigilance, pattern recognition, and microsecond reaction times.

In practice, this means:

- AI hunts, neutralizes, and quarantines threats automatically.
- Human teams review the AI's actions, fine-tune policy, and investigate anomalies beyond the AI's context.
- Continuous feedback loops ensure that the AI evolves in alignment with organizational goals and ethical boundaries.

Challenges and Dark Sides

The leap to autonomous AI defense isn't without risk:

- **Overreach & False Positives:** An overly aggressive AI could misclassify legitimate traffic as malicious, causing costly downtime.
- **AI vs. AI Battles:** Offensive AIs are already in development — future breaches may be waged entirely between machine agents with humans sidelined.
- **Explainability:** Fully autonomous systems must be able to justify their decisions to regulators, auditors, and courts — no “black box” defense will be acceptable.

And then there's the inevitable **arms race**: once defensive AI becomes predictive and autonomous, offensive AI will adapt in kind, leading to real-time, continuous cyber battles fought entirely without human intervention.

Why Now? The Convergence of Capabilities

Several technological shifts are aligning to make this possible today:

1. **Transformer-based AI Models** – Capable of reasoning about system state changes in ways traditional ML models never could.
2. **Edge AI Processing** – Puts autonomous defense capabilities closer to the attack surface, reducing latency to near-zero.
3. **Global Threat Intelligence Integration** – Merging real-time feeds with AI analysis creates a dynamic, self-updating defense posture.

This convergence means the leap to autonomous predictive defense isn't decades away — it's already underway in classified government projects, high-security financial institutions, and bleeding-edge tech companies.

The Strategic Imperative

For CISOs, the question is no longer *if* they will integrate autonomous AI into their security stack, but *how quickly*. The early adopters will gain:

- **Massive Reduction in Dwell Time** – Intrusions can be identified and neutralized in seconds rather than weeks.
- **Operational Cost Savings** – AI handles repetitive, high-volume triage work.
- **Competitive Advantage** – Secure infrastructure is now a business differentiator, especially for industries handling sensitive customer data.

Those who wait risk facing adversaries armed with AI that can dismantle their defenses faster than their teams can open a ticket.

Looking Ahead: AI as the Immune System of the Internet

If fully realized, autonomous defensive AI could evolve into a **planetary-scale immune system** — distributed, adaptive, and constantly learning. Every attack detected on one node would instantly inoculate every connected system worldwide.

This vision, however, demands unprecedented collaboration between governments, private sector companies, and open-source communities — a level of trust and shared mission that cybersecurity has historically struggled to achieve.

Final Warning: The Ghost Is Already Here

The term “ghost in the firewall” isn’t just poetic. The first truly autonomous defensive AIs will operate invisibly, hidden in the noise of network telemetry, watching and countering threats before anyone notices.

When it works, the effect will feel almost supernatural — breaches that never happen, exploits that mysteriously fail, and adversaries who can’t explain why their payloads vanish into the void.

The cybersecurity world has always talked about “staying one step ahead of the attacker.” In the age of autonomous AI, that step becomes **an infinite lead** — but only for those willing to let the ghost in.


About the Author

Bekir Tolga TUTUNCUOGLU is the CEO of TTNC Teknoloji. He is an internationally recognized expert in artificial intelligence and cybersecurity, with over 15 years of pioneering research, enterprise consulting, and innovation in the field. As a keynote speaker at global technology conferences, a published thought leader in leading industry journals, and a fellow member of The Institution of Engineering and Technology, he has shaped forward-looking strategies for institutions, international companies, and emerging tech startups.

His work spans designing advanced AI-driven threat detection systems, advising on national cyber defense frameworks, and leading high-impact initiatives that bridge cutting-edge academic research with real-world applications. He has served as a mentor and reviewer for global innovation challenges, and his contributions have been cited in top-tier media and policy discussions.

Bekir Tolga TUTUNCUOGLU can be reached online at <https://www.linkedin.com/in/tolga-tutuncuoglu/> and at our company website <https://hozzt.com>





The New Face of Cyber Threats: From Scattered Spider to North Korea's Phantom Workforce

And Why Cyber Threat Intelligence Is Our Best Defense

By Richard K. LaTulip – A Field Chief Information Officer at Recorded Future

In today's hyperconnected digital world, the threat landscape has evolved far beyond lone hackers and opportunistic malware. Cyber adversaries have become more structured, strategic, and in some cases, state-sponsored. Two examples have dominated recent headlines: Scattered Spider, a loosely affiliated yet highly effective cybercriminal syndicate, and North Korea's phantom information technology (IT) workforce, a clandestine network of fake employees infiltrating real companies under false pretenses. While these actors differ in motivation, tactics, and structure, they share one critical commonality: they thrive in environments lacking proactive cyber threat intelligence (CTI) integration.

This article explores how these threat groups operate, the vulnerabilities they exploit, and most importantly, why organizations must embed CTI into the foundation of their strategic defense plans, not just as a technical feed but as a driver of risk-aware decisions across the business.

Scattered Spider: The Rise of Decentralized, Agile Threat Actors

Scattered Spider, also tracked as UNC3944 or under aliases such as Octo Tempest, is not your average ransomware gang. Emerging in 2022 and quickly making headlines with high-profile attacks against major enterprises like MGM Resorts and Caesars Entertainment, the group has become emblematic of a new breed of threat actor. According to public [reporting](#), it took the attackers less than **ten minutes** to disrupt several critical systems inside MGM's enterprise an astonishing demonstration of speed and precision.

So who are they? Current intelligence suggests that Scattered Spider is composed primarily of young, English-speaking threat actors who are highly skilled in social engineering. This was recently corroborated by [KerbsonSecurity](#), which reported that the UK's National Crime Agency (NCA) arrested four individuals between the ages of 17 and 20 in connection with the group.

But it's not their age that makes Scattered Spider particularly dangerous; it's their agility. Unlike traditional threat actors that rely heavily on malware or zero-day exploits, this group exploits human error, trust, and identity to gain access. Their tactics include:

- Vishing and phishing campaigns targeting help desk personnel
- SIM swapping to intercept multi-factor authentication (MFA) tokens
- MFA fatigue attacks—bombarding users with push notifications until one is accepted
- Exploitation of identity platforms such as Okta and Active Directory

These methods enable rapid lateral movement, often culminating in ransomware deployment once persistence is achieved. Their operational ties to the ALPHV (BlackCat) ransomware group mark an evolution from initial access brokers to full-spectrum threat actors.

So, how could cyber threat intelligence have made a difference?

CTI empowers organizations to track threat actor tactics, techniques, and procedures (TTPs) in near real-time. When integrated into detection engineering and incident response workflows, CTI helps anticipate social engineering ploys, enrich security alerts with adversary context, and inform identity-centric risk scoring. Behavioral intelligence fueled by CTI can detect early-stage anomalies such as unexpected SIM swaps or login patterns long before ransomware detonation.

Moreover, CTI enables defenders to shift from reacting to breaches to proactively hunting for precursor domains, IPs, toolsets, and infrastructure associated with known threat groups like Scattered Spider. Essentially, CTI turns the unknown into the observable, which is crucial for rapid defense.

North Korea's Phantom Workforce: State-Sponsored Insider Threats

If Scattered Spider is a masterclass in decentralized criminal agility, North Korea's IT worker program is a calculated, state-backed strategy that exploits trust from the inside. In recent years, the U.S.

government, through joint bulletins from the FBI, CISA, and Treasury, has warned about thousands of North Korean nationals posing as freelance developers to obtain remote jobs at Western tech companies.

In one such [advisory](#), the FBI highlights how North Korean IT workers leverage U.S.-based individuals, both witting and unwitting, to gain fraudulent employment. Using forged documents, stolen identities, and clean GitHub profiles, these workers are paid in cryptocurrency and often routed through intermediaries to obfuscate their origins. Ostensibly harmless, many perform routine software development tasks. But the risk isn't hypothetical; it's geopolitical:

- These roles generate foreign currency for North Korea's weapons programs.
- Some workers are suspected of introducing backdoors or exploitable code into legitimate software.
- Others have engaged in IP theft, data exfiltration, or indirect supply chain compromise.

As [the FBI warns](#), North Korean IT workers have been “leveraging unlawful access to company networks to exfiltrate proprietary and sensitive data, facilitate cyber-criminal activities, and conduct revenue-generating activity on behalf of the regime.”

Unlike traditional external breaches, these “hires” are **inside the perimeter**, often with privileged access to repositories, internal tools, and CI/CD pipelines.

This is where CTI is essential not just for security teams but also HR, legal, and compliance functions. For instance:

- CTI platforms can monitor for aliases tied to DPRK entities, including reused email addresses, wallets, or infrastructure.
- Geo-behavioral analytics can flag inconsistencies in time zones, login locations, or regional language usage.
- Integration with sanctions intelligence can alert companies to prohibited payments or crypto flows.

Without CTI, even the most advanced EDR tools won't detect a North Korean contractor silently pushing a Git commit from a coffee shop in Serbia. With CTI, organizations gain contextual insight into digital identities, allowing proactive responses to nation-state insider threats.

From Intelligence Deficiency to Strategic Integration: Making CTI Work for the Enterprise

Despite their divergent models, Scattered Spider's chaotic ingenuity and North Korea's methodical deception both adversaries exploit the same underlying weakness: the absence of integrated cyber threat intelligence across the enterprise. These actors don't merely compromise systems; they take advantage of organizational silos, outdated defenses, and fragmented approaches to risk management.

As noted in recent [analysis](#), the digital ecosystem is more interconnected and exposed than ever. With the rise of AI-driven attacks, supply chain compromises, and zero-day vulnerabilities, cyber threat intelligence has become essential to proactive defense. Yet many organizations still view CTI as merely a technical tool focused on a stream of indicators of compromise stored in the security operations center

(SOC) and loosely connected to alert triage. This approach fails to account for how modern adversaries operate: not through easily detectable artifacts, but through social engineering, identity compromise, and behavioral manipulation.

Static rules and compliance checklists are no match for such threats. As Levi Gundert, Chief Security and Intelligence Officer at Recorded Future, explains, "Threat intelligence is essential to modern cyber-risk management and resilience for two reasons. First, AI-enabled adversaries move at machine speed, so intelligence must flow automatically into digital-risk, cyber-operations, exposure-management, and control-validation systems. Second, although compliance often drives security budgets, executives need an intelligence advantage to invest in the right controls at the right time."

To remain effective, CTI must evolve beyond a reactive feed and become a strategic risk function one that connects intelligence, operations, and executive leadership. That requires:

- Informing procurement and hiring processes through intelligence-backed vetting
- Enabling identity-aware detection strategies across infrastructure
- Conducting red team exercises and tabletop simulations based on known threat actor playbooks

When CTI is underutilized, organizations are flying blind. But when it is integrated into the business, it becomes the lens through which complex threats are identified, evaluated, and mitigated before they cause harm.

Acknowledging this intelligence gap is only the first step. Organizations must now operationalize CTI as a cross-functional capability that aligns with business priorities, technical controls, and governance frameworks.

To achieve that, CTI must become a strategic pillar, deeply integrated into the enterprise's cybersecurity architecture and decision-making process. This includes:

- Identity-Centric Defense: Embedding CTI into identity and access management systems to flag high-risk authentications and enrich access logs with threat context
- Security Operations Integration: Feeding CTI into SIEM and SOAR platforms to improve alert triage, risk-based prioritization, and SOC efficiency
- Proactive Threat Hunting: Using CTI to build behavioral detection aligned with specific adversary tactics, such as MFA fatigue or suspicious code pushes
- Board-Level Risk Visibility: Translating intelligence into risk narratives that resonate with executives and board members, answering not just what happened, but what it means

CTI is no longer the domain of threat analysts alone. It is a business-wide capability that informs strategic planning, drives smarter investment decisions, and enhances organizational resilience. It empowers CISOs to allocate budgets effectively, enables legal teams to assess third-party exposure, and supports developers in writing secure code grounded in real-world adversary behavior.

In today's threat landscape, connecting the dots isn't optional anymore. Cyber threat intelligence provides the context, foresight, and speed required to see the full picture before the adversary makes their next move.

Conclusion: Adapt or Be Compromised

Cyber adversaries are evolving and fast. Whether it is Scattered Spider exploiting identity-based trust through sophisticated impersonation, or a North Korean IT worker silently embedding risk within your software pipeline, the modern threat landscape demands more than traditional perimeter defenses. It demands intelligence.


Cyber threat intelligence is not a luxury. It is a strategic necessity. It empowers organizations to anticipate threats, adapt in real time, and respond with precision. As adversaries become faster, more deceptive, and more embedded, CTI must serve as the connective tissue across operations, governance, and risk management.

Understanding the adversary is no longer a reactive exercise. It is the foundation of modern cyber defense. And that foundation begins with intelligence. End of article.

About the Author

Richard K. LaTulip is the Field Chief Information Officer at Recorded Future. He advises global organizations on cyber threat intelligence strategy, risk management, and security operations. He is a retired U.S. Secret Service agent with over two decades of experience investigating cybercrime, leading long-term undercover operations, and collaborating with international law enforcement. Richard holds a Master of Science in Cybersecurity Policy and Governance from Boston College and maintains several industry certifications, including CISSP, CISM, CEH, CySA+, C|CISO, and Security+. He is also the author of *Operation Carder Kaos: How One Agent Penetrated the Underground Community*, a forthcoming book chronicling his infiltration of cybercriminal networks. He regularly speaks on topics related to cyber threats, resilience, and intelligence-driven defense. Rich can be reached online at Richard.latulip@recordedfuture.com or [linkedin.com/in/richard-latulip-5852606](https://www.linkedin.com/in/richard-latulip-5852606) and at our company website <https://www.recordedfuture.com/>.





The Real Cost of Exposure Remediation: Helping Developers Avoid Burnout

By Ravid Circus, Chief Product Officer, Seemplicity

The number of security weaknesses companies face keeps growing fast. Common vulnerabilities and exposures (CVEs) jumped by 39% in 2024, adding to the overwhelming burden on remediation teams. Whether these risks are found in owned code, third-party code or commercial software, they don't just compromise organizational security; these risks contribute to the ongoing cybersecurity workforce crisis, with developers, in particular, struggling to keep pace with fixing the never-ending backlog of vulnerabilities.

Currently, 69% of those contributing to remediation are developers, highlighting that, on top of their core responsibilities, developers play a critical role in fixing vulnerabilities. With the added workload, developers are at greater risk of experiencing burnout, and because burnt-out workers are more likely to leave their jobs, organizations risk seeing declining productivity and higher employee turnover unless this workload is alleviated.

The burden of remediating a vast number of vulnerabilities can also have an impact on organizations' security resilience. Teams that are stretched too thin are more likely to take a scattered, non-strategic, reactive approach to remediation, causing them to overlook critical vulnerabilities and delay fixes. This approach leaves organizations exposed and can lead to harmful outcomes, including more security incidents, financial losses, and an erosion of stakeholder and customer trust.

For organizations to maintain a strong security posture and set their teams up for success, they must adopt proactive exposure management strategies. To do this, organizations should implement a plan that prioritizes vulnerabilities based on business impact, streamlines workflows, and improves collaboration among fixing teams.

Prioritize Vulnerabilities Based on Business Impact

Many security organizations focus their efforts on identifying vulnerabilities, using an [average of eight security scanning tools](#) to do so. Since these tools provide limited context and visibility into risks, they often just add to the noise that security teams need to manage. This can cause teams to focus on less critical vulnerabilities while more significant ones remain unaddressed.

Only focusing on identifying vulnerabilities causes organizations not to be able to differentiate between low-impact vulnerabilities and high-impact vulnerabilities that require immediate action. Beyond missing important vulnerabilities, using non-strategic, manual methods based on guesswork to fix the low-impact issues wastes valuable time and resources that could be spent on remediating critical vulnerabilities.

Implementing a risk-based exposure management strategy using AI-driven technology can enable organizations to analyze vast amounts of data and prioritize vulnerabilities that pose the greatest business risk. A more strategic approach helps provide relief for fixing by zeroing in on pressing issues.

Streamline Workflows with Automation

Manual processes for vulnerability remediation lead to significant inefficiencies. Developers and other teams contributing to remediation efforts waste valuable time manually identifying the most critical vulnerabilities, assigning remediation steps to the appropriate teams and coordinating the fixes. Particularly when there's a large number of vulnerabilities involved, this approach can quickly become overwhelming.

According to Seemplicity's [2024 Remediation Operations Report](#), nearly 60% of organizations still incorporate manual processes in their remediation efforts. Relying too much on individual contributors can disrupt workflow consistency and efficiency. Given that many security teams operate with constrained resources and limited staff, it's almost impossible to keep up with the volume of vulnerabilities that need to be addressed.

Automating processes and workflows doesn't just enhance efficiency - it empowers teams by reducing manual work, boosting productivity and ensuring that time and energy is focused on efforts that make a difference to organizations' security resilience.

Cross-Team Collaboration Across Remediation Teams

Effective exposure management is a shared responsibility that requires collaboration between security and the teams doing the remediation, but in many organizations, teams operate in silos without established communication channels.

Lack of coordination, poor communication and unclear ownership of tasks can cause internal friction and inefficiencies, which negatively impact team mental health, exacerbate workload issues and create remediation gaps that can increase an organization's threat exposure. To foster stronger collaboration, organizations should integrate remediation workflows and tasks into the existing processes and tools developers use. This enables teams to address vulnerabilities quickly and efficiently while creating a healthier, sustainable work environment.


Support Fixing Teams to Strengthen Security

The success of any exposure management strategy depends on the well-being of the teams responsible for executing it. The human impact on security operations cannot be overlooked: Developer, operations and security teams are all critical players when it comes to protecting organizations from vulnerabilities.

Security leaders need to stop simply reacting to problems as they come up and start getting ahead of them. They need an exposure management strategy that makes the company more secure and protects their most valuable asset: their people.

About the Author

Ravid Circus is Co-founder and Chief Product Officer at Seemplicity. Ravid has 20+ years of experience in translating risk management processes to technology. As former VP of Customer Success and VP Products at Skybox Security, Ravid has a unique perspective on cyber security management and hands-on experience with the technology required to support it. As a security practitioner at heart, Ravid understands the customer's technology and operational challenges around risk reduction. His years of deploying customer care initiatives at Credit Suisse, Citi and Chase earned him a reputation as a seasoned security technologist. Ravid can be reached on [LinkedIn](#) and at our company website <https://seemplicity.io/>





The Rise of AI-Driven Credential Stuffing: Why IAM Alone Can't Save You

When Bots Learn to Think Like Humans

By Sandeep Dommari, Principal Architect, Ping Identity

Introduction: When Bots Learn to Think Like Humans

Security teams dismissed credential stuffing as "noisy bot traffic" in 2012. It has now developed into one of the world's most lucrative, scalable, and AI-powered threats.

Consider the recent spate of hacks at Nintendo, Zoom, and Spotify. Attackers only needed to use usernames and passwords that had been stolen from unrelated leaks and allow bots to test them across millions of accounts; they didn't even need to breach the companies directly.

Attackers no longer simply spray stolen credentials thanks to generative AI. They are so good at imitating human behavior that they are fooling even sophisticated Identity and Access Management (IAM) systems and conventional bot detection tools.

The Evolution of Credential Stuffing Beyond "Just Bots"

In the past, credential stuffing involved using speed bots to brute force login forms with compromised usernames and passwords.

- They were frequently prevented by security measures like IP blacklisting, velocity checks, and CAPTCHA.
- The New World with AI: AI-powered bots dynamically modify attack velocity, device fingerprints, and session patterns through reinforcement learning.
- Attackers repurpose tools such as open-source machine learning libraries to teach bots to "look" like actual users.
- Credential stuffing campaigns now circumvent MFA by taking advantage of push notification fatigue or weak SMS-based factors (also known as "MFA bombing").

Real-world illustration:

An AI-driven attack against a multinational retail behemoth in 2023 involved bots that mimicked human shopping behavior by rotating IP addresses across mobile networks, simulating device orientation sensors, and even varying login attempts over several weeks. There was nothing unusual in the IAM logs. The fraud bill? Chargebacks and account takeovers totaling more than \$20 million.

The Business Impact: The Importance for CEOs and Boards

Credential stuffing is more than just an IT annoyance to a CEO or board member. There are repercussions from this business continuity risk:

- **Brand Damage:** Even if credentials originated from another breach, the victim always holds the company accountable when a user account is compromised.
- **Fraud Costs:** Mostly as a result of credential stuffing, airlines report \$750 million in loyalty program fraud each year.
- **Operational Disruption:** Calls for password resets and fraud investigations overwhelm customer support teams.
- **Regulatory Penalties:** In accordance with the GDPR/CCPA, there may be fines for neglecting to protect customer accounts, even from credentials that are reused.

For instance, Nintendo acknowledged in 2020 that a credential stuffing campaign had compromised 160,000 accounts. Parents and regulators were outraged when the attackers targeted children's accounts for stored credit card data.

Why You Can't Be Saved by Me Alone

IAM systems are essential, but they are not made to handle this issue.

- Although attackers can use AI to evade MFA prompts, IAM can enforce strong authentication.
- Although IAM can centralize identities, APIs and shadow apps introduce vulnerabilities.
- Although AI-driven bots produce noise that appears to be human logs, IAM offers audit logs.

IAM is reactive, which is a painful reality. Layered defenses and proactive detection are necessary for credential stuffing.

How AI Enhances Credential Stuffing

Attackers use AI as a weapon in the following ways:

1. **Behavioral Mimicry:** By recording actual user sessions, bots teach machine learning models to imitate mouse motions, geolocation switching, and typing rhythm.
2. **Learning that Adapts:** Bots use residential proxies to route if IPs are blocked; they learn from unsuccessful login attempts. They target distinct accounts if MFA is activated.
3. **Changes to Passwords:** Passwords that have been stolen are transformed into dozens of different variations by generative models ("Summer2023!" → "\$umm3r2023!!").
4. **Getting around CAPTCHAs:** The majority of CAPTCHAs are broken at scale by computer vision and LLM-powered solvers.
5. **Abuse of APIs:** Bots circumvent web defenses by directly exploiting login APIs.

Case Study: In order to evade fraud detection, a financial services company found that attackers were utilizing reinforcement learning bots that modified login attempts in real time. Before the attack was lessened, it took six months and a new bot defense solution.

Developing a Defense Outside of IAM

1. Feeds of Credential Intelligence

- Incorporate threat intelligence that keeps an eye on dark web dumps and sends out alerts when user credentials show up.
- For instance, businesses that use these feeds proactively reset exposed accounts following the LinkedIn leak.

2. AI-Powered Bot Detection

- Use anomaly detection at the edge (WAF/CDN), which examines behavioral patterns such as device fingerprinting, velocity, and mouse movement entropy.
- Best practice: Integrate in-house tuning with vendor solutions (like PerimeterX and Akamai Bot Manager).

3. Authentication that Adapts

- Get rid of static MFA. Make use of risk-based policies:

- Only use step-up authentication in cases where the session risk score is high.
- Take into account device history, IP reputation, geolocation, and impractical travel.

4. Security Controls for APIs

- Keep web login flows and login APIs separate.
- As an illustration, apply rate limits and schema validation, especially for mobile app login APIs.

5. Transparency & User Education

- Inform clients about the dangers of password reuse and send out password breach notifications when credentials are found to be reused.
- Openness fosters trust.

A Guide for CEOs and CISOs

Phase 1: Visibility & Discovery

- Map every login endpoint, including mobile flows and APIs.
- Benchmark failed login ratios: suspicious increases frequently indicate credential stuffing.

Phase 2: Incorporate Controls

- Enhance the current IAM with threat intelligence, bot detection, and adaptive authentication.
- Enhance IAM rather than completely replace it.

Phase 3: Ongoing Examination

- Conduct red-team drills that replicate credential stuffing in particular.
- Apply the concepts of chaos engineering to test the resilience of your login flow in the event of an attack.

Phase 4: Involvement with Businesses

- Report credential stuffing as revenue loss and customer fraud rather than a "login failure issue."
- Frame board talks about revenue and reputational risk.

Conclusion: The Arms Race in AI Has Started

Automation versus defense has always been at the center of credential stuffing. However, AI has ushered in a new era of intelligent, adaptable, and outsmarting robots.

The lesson for CISOs and CEOs is straightforward: IAM is important, but not enough.

Those who invest in layered defenses, test resilience frequently, and treat credential stuffing as a systemic business risk will emerge victorious.


If not, attackers will be able to walk through your digital front door rather than just knock on it.

About the Author

Sandeep Dommari is a Senior Cybersecurity Architect and IAM Strategist with over 18 years of experience designing secure access frameworks across Fortune 100 enterprises. His work focuses on application security, adaptive identity, and building secure-by-design architectures for critical industries.

Sandeep can be reached online at sandeep.dommari@ieee.org





The Role of Agentic AI in Proactive Cyber Threat Hunting

By Nivedita Kumari, Data & AI Customer Engineer, Google

As cyber threats continue to evolve and grow more persistent and sophisticated, simply applying tools and technologies to respond to previous attacks is no longer enough. Organizations have to move away from defensive strategies and proactively define threats that could already be lurking unnoticed in their environments. In this context, agentic AI is being identified as a transformative tool for the practice of cyber threat hunting.

Understanding Agentic AI in Cybersecurity

Agentic AI refers to systems that enable artificial intelligence (AI) to autonomously observe their environment, act, and make decisions to achieve a predetermined goal. Agentic AI can be deployed within networks and systems, permitting the AI agents to monitor activities in order to analyze data and actively hunt for dormant or emerging threats without the need for sustained human oversight.

Agentic AI is the first form of artificial intelligence that is not bound to rule-based systems or supervised models that are defined by pre-established heuristics and labeled examples. They have the ability to learn and adapt to new threats in real time and are able to detect nuanced anomalies and indicators of compromise (IOCs) that can be overlooked by other security tools or human experts.

How Agentic AI Enhances Proactive Threat Hunting

Agentic AI has some significant benefits for proactive cyber threat hunting:

- **Autonomous Threat Discovery:** While an organization understands the systems, networks, and devices it depends on, an agent arm can independently explore through the data sources, network traffic, system logs, and endpoint activity looking for ignored threats that others have missed in the various security layers.
- **Identification of Dormant Threats:** Cyber threat hunting agents are able to discover threats that are not actively causing damage, but they may trigger if provoked. An agent can accumulate and analyze the behavioral baseline over time, allowing detection of anomalies that may be an initial signal of future malicious activity.
- **Detection of Emerging Threats:** Agentic AI can also learn from what it is exposed to and can change with changes in attacks. If a system is attacked with a new technique and relevant training data are available, the agent could detect what others cannot, opening up new avenues of detection for zero-day exploits and novel malware.
- **Reduced Alert Fatigue:** As AI agents are triaging and investigating the findings autonomously, they will discourage the churning of alerts most organizations experience which are predominantly false positives or low-value alerts. This will allow analysts to focus on the important issues they need to address and other contextual incidents that warrant their attention.
- **Continuous Monitoring and Analysis:** AI agents are a 24/7 system of monitoring and analysis of the cyber security environment to identify any threats are being dealt with appropriately.
- **Enhanced Threat Intelligence:** The insights & findings generated by agentic AI can help advance understanding of the threat landscape by vastly improving threat intelligence and informing future security responses.


The Future of Threat Hunting with Agentic AI

Agentic AI is not intended to fully replace human threat hunters. It is designed to complement human capabilities. With repetitive, time-consuming, and resource-heavy aspects of threat hunting automated by an AI agent, a human analyst will have more time for strategic thinking, complex analysis, and incident response.

Agentic AI is progressing toward the future where the defense of our cyber systems becomes preemptive, adaptive, and resilient against threats that are always evolving. With advancing capabilities, AI agents will be an increasingly important asset in the protection of digital assets in the face of cyber threats. Organizations that use agentic AI to perform proactive threat hunting will have a distinct advantage in the protection of their digital assets against their adversaries.

About the Author

Nivedita Kumari is a Data & AI Customer Engineer at Google, Nivedita is dedicated to making machine learning and AI accessible to everyone. She works to break down the complexities of these technologies, empowering individuals from all backgrounds to explore and understand AI's potential. Her focus is on driving tangible business value through data-driven strategies, while also inspiring and mentoring the next generation of AI innovators. Nivedita can be reached at <https://www.linkedin.com/in/nivedita-kumari/> and at our company website <https://www.google.com/>.





The Security-speed Myth That's Sabotaging Your Modernization

Why The Biggest Obstacle to Legacy Upgrades Isn't Technology—it's How We Organize IT Decisions

By Ayo Akinsanya, CISSP, PMP, ITIL (Cybersecurity Expert)

The Modernization Standoff

Legacy system upgrades consistently follow the same pattern: leadership demands rapid deployment to meet business objectives, while security teams require extensive planning to address compliance and risk concerns. The result is often organizational paralysis, with projects delayed indefinitely or rushed to market with dangerous vulnerabilities.

This conflict stems from a deeper organizational problem than most realize. It's not just about choosing between security and speed - it's about how we structure IT decision-making, allocate budgets, and

measure success. The security-versus-speed dilemma is often a symptom of misaligned incentives rather than technical limitations.

While some security-speed conflicts are genuine and unavoidable, many organizations create false choices through outdated governance structures, risk-averse cultures, and siloed decision-making processes.

Why Smart Organizations Keep Making Dumb Choices

The persistence of security-speed conflicts reveals fundamental flaws in how enterprises approach technological decisions:

Budget Structure Problems: Security and operations have separate budgets, creating artificial competition for resources rather than collaborative investment in shared outcomes.

Risk Culture Misalignment: Security teams are penalized for incidents but not rewarded for enabling business velocity, while operations teams face pressure for speed but limited accountability for security outcomes.

KPI Disconnection: Success metrics rarely account for interdependence between security and operational efficiency, leading to suboptimal decisions that look good on individual scorecards.

These flawed structures persist because changing them requires short-term pain for long-term gain. CFOs resist unified budgets because they complicate financial tracking. Security leaders fear losing autonomy over risk decisions. Operations teams worry about accountability for incidents they can't control.

But here's the uncomfortable question: Are these structural problems, or do they reflect genuine technical realities?

The Adaptive Security Counter-Argument

Some security professionals argue that speed pressure inevitably leads to dangerous shortcuts. They contend that adaptive approaches are sophisticated ways to rationalize compromised security.

The Skeptical View: "Context-aware security sounds appealing, but real-world implementation means someone - usually business stakeholders - decides what constitutes 'acceptable risk.' This invariably leads to gradual erosion of security standards under business pressure."

The Technical Reality: Adaptive security implementations require significant upfront investment in infrastructure, tooling, and expertise. Organizations lacking technical maturity may implement these approaches poorly, creating an illusion of security while introducing new vulnerabilities.

The Genuine Conflicts: Certain scenarios present unavoidable tradeoffs. Regulatory requirements sometimes mandate specific controls regardless of operational impact. Emergency situations may require accepting risk to maintain critical services.

When Adaptive Approaches Actually Work

Progressive organizations are successfully applying adaptive security principles, but success depends on specific conditions:

Technical Prerequisites

Infrastructure Maturity: Organizations need API-driven architectures with micro segmentation capabilities, real-time policy engines, and automated compliance validation - typically requiring cloud-native or hybrid infrastructures less than five years old.

Data Visibility: Adaptive approaches require user behavior analytics platforms, network traffic analysis, endpoint detection and response tools, and security orchestration platforms that can process and correlate data in under 60 seconds.

Automation Capabilities: Success requires automated policy enforcement and response capabilities.

Organizational Prerequisites

Unified Success Metrics: Teams must share accountability for both security and operational outcomes rather than optimizing for conflicting KPIs.

Risk-Informed Culture: Organizations need sophisticated risk assessment capabilities and cultures that can make nuanced security decisions rather than defaulting to maximum protection.

Technical Leadership: Implementation requires deep expertise in both security and operations, not just project management skills.

The Broader Implications

If security can genuinely enhance rather than hinder operations, what does this reveal about traditional IT organizational structures?

Governance Evolution: Organizations may need to restructure IT governance to optimize for combined security-operational outcomes rather than managing them as separate functions.

Role Redefinition: Traditional security roles focused on risk mitigation may evolve toward "operational enablement through security intelligence."

Investment Philosophy: IT budgeting might shift from cost center thinking toward strategic capability investment that delivers both security and business value.

The Uncomfortable Questions

Resource Reality: Adaptive security often requires more sophisticated infrastructure and expertise than traditional approaches. Are organizations prepared for these investments?

Cultural Readiness: Many organizations claim to want innovation but revert to risk-averse behaviors when implementation challenges arise. How many are genuinely prepared for the cultural changes adaptive approaches require?

Technical Honesty: Are adaptive security capabilities hyped, or do these approaches genuinely resolve fundamental conflicts between security and operational efficiency?

Implementation Reality Check

The Resource Paradox: Organizations most needing adaptive security - those with legacy systems and security-speed conflicts - often lack the infrastructure, expertise, and agility for successful implementation. Early adopters compound advantages while laggards face increasing costs.

Prerequisites for Success

- Modern, API-driven architectures with comprehensive monitoring
- Teams that already collaborate across security-operations boundaries
- Sustained investment commitment beyond initial project funding

Warning Signs to Reconsider

- Legacy systems with inflexible security models
- Highly prescriptive regulatory environments
- Organizations struggling with basic security hygiene

The Strategic Decision

The security-versus-speed debate reflects deeper questions about organizational maturity, risk management philosophy, and technical architecture decisions. While adaptive security approaches can resolve many false choice scenarios, they require substantial organizational commitment and technical sophistication.

The critical question isn't whether adaptive security works—it's whether your organization has the prerequisites for successful implementation. Honest assessment of technical capabilities, cultural readiness, and resource availability is essential before pursuing these approaches.

Organizations that can successfully implement adaptive security gain significant competitive advantages. Those that attempt implementation without proper prerequisites may create new problems while failing to solve existing ones.

About the Author

Ayokunle (Ayo) Akinsanya is a cybersecurity and secure enterprise implementation expert with over 16 years of experience, recognized for developing revolutionary frameworks that resolve the traditional conflict between security and operational efficiency.

His practical implementations have protected millions across critical infrastructure sectors during complex enterprise deployments.

Akinsanya holds a master's in information systems & sciences from Bowie State University and premier certifications including CISSP, PMP, CBAP, and ITIL. He serves on the National CyberWatch Center Curriculum Standards Panel and has published strategic methodologies in ISC2 Insights, the premier cybersecurity professional publication certified professionals globally.

His Adaptive Security Architecture frameworks continue influencing cybersecurity professionals worldwide, strengthening global defenses while enabling operational excellence across healthcare, manufacturing, and financial sectors.

Ayo can be reached online at kunlay003@gmail.com, LinkedIn <https://www.linkedin.com/in/ayokunle-akinsanya-1402/>





The Top Three Emerging AI Threat Tactics Changing the Face of Identity Fraud

By Ashwin Sugavanam, VP of AI & Identity Analytics

AI innovation is accelerating at a pace the world has never seen before and unfortunately providing just as many advancements for bad actors as it is for good. It is now alarmingly easy to manufacture believable, fake digital identities, and these advancements are outpacing traditional identity verification methods. For example, initial biometrics-based identity verification tactics once offered sufficient barriers, but today's cybercriminals are using tools that can manipulate even sophisticated security systems.

From injected selfies to synthetic identities, these fraud techniques are industrialized, automated, and scalable. If we don't evolve our defenses to adapt to these threats, we risk opening the door to widespread exploitation.

Public Trust Is Eroding

For enterprise security leaders, it's important to note just how much of an impact deepfakes have had on eroding the public's trust. In a recent global study, [69% of consumers said AI-powered fraud](#) now poses a greater threat to their personal security than traditional forms of identity theft. As deepfakes and synthetic media grow more convincing, the line between real and fake continues to blur, and with it, confidence in online authenticity.

When asked who should be responsible for stopping these threats, 43% of consumers pointed to Big Tech, while just 18% believed the responsibility lies with themselves. Users have a growing expectation for enterprises to keep their digital identities safe while engaging with their platform. As manipulated content becomes harder to detect, service providers must adapt their identity intelligence systems to meet the growing needs of their customers.

So, what exactly are today's most pressing identity fraud threats, and how can AI help us detect and stop them? Let's break down three of the most concerning threats on the horizon, and the AI-powered identity intelligence strategies that can stop them in their tracks.

1. Injected Selfies & Deepfakes: Breaking Biometrics

Facial recognition is the foundation of biometric authentication. But what happens when the face isn't real?

Today, cybercriminals are increasingly turning to camera injection attacks, which are digitally inserted images that bypass the physical camera feed. These attacks, sometimes carried out through virtual cameras, enable fraudsters to simulate a live video stream using pre-rendered or AI-generated faces. Camera injection attacks are the delivery mechanism behind most deepfake-based fraud, but they're also being used as standalone attack vectors.

Deepfake fraud, along with wreaking havoc on social media platforms and tarnishing public figures' reputations, is now impacting executives in the boardroom. Hackers are leveraging advanced deepfake tools to hijack high-level executives' identities and exploit the enterprise. All it takes is a few images, a few seconds of a digitally recorded voice sample, and an unsuspecting employee to be scammed into performing costly actions by the fraudster.

One answer to stopping these attacks is the use of liveness detection. This strategy uses AI-based algorithms that assess the authenticity of the person behind the camera in real time. Advanced systems can detect micro-movements, light reflections, and other physiological cues that can't be faked. Multimodal liveness checks, which combine visual, auditory, and motion-based signals, are quickly becoming the gold standard.

2. Synthetic Identities: A Quiet but Devastating Attack Vector

In the arms race of digital fraud, synthetic identity creation has become one of the most challenging threats to detect. Fraudsters build completely fake identities using real data fragments, such as legitimate Social Security numbers, and couple that information with forged addresses to replicate a new identity across multiple accounts.

As document fraud becomes increasingly sophisticated, relying on static template matching is no longer viable. Security teams must adopt document liveness models that can detect PII data manipulation. This includes the absence of security features like holograms, photocopies, digital copies, and superimposition of the photo area to establish document validity.

3. Uncovering organised fraud rings

A pattern we observe typically with fraud rings is that ID & Selfie images in most cases have similar backgrounds. This visual uniformity is often overlooked by traditional ID verification systems, allowing multiple fraudulent identities to slip through as seemingly legitimate.

Networked AI models, which evaluate the broader behavior of user identities across identity transactions, devices, and IPs, can help detect patterns indicative of synthetic identity networks. By combining cross-transactional risk signals, identity intelligence and document forensics, organizations can uncover clusters of fraud that may otherwise go unnoticed. This includes detecting both organized fraud rings operating across multiple digital platforms and individual users attempting scams on multiple fronts.

Fighting Fraud with AI

Traditional verification methods, even when automated, can't stand alone against today's AI-enhanced threats. Security teams must adopt a layered approach powered by advanced identity intelligence.


This means embracing biometrics-based security tactics like liveness detection to combat camera injection and deepfakes, and incorporating adaptive, cross-transactional risk signals to expose synthetic identity networks.

With AI working on our side, not just against us, we can build adaptive systems that meet fraudsters at their level and stay one step ahead.

About the Author

Ashwin Sugavanam is currently the VP, AI & Identity Analytics at Jumio Corporation is a visionary Data and Analytics leader with two decades of overall experience out of which he has spent the last decade in helping organizations incubate and scale Data & AI practices.

Over the last couple of years Ashwin has helped organizations drive measurable business outcomes by responsibly scaling Data and AI initiatives, implementing modern concepts like Data Mesh and MLOps, and leveraging tools such as the Data Scientist Co-Pilot to accelerate impact. He can be reached on [LinkedIn](#) and at our company website <https://www.jumio.com/>.





The AI Arms Race Is Here: How Enterprises Must Weaponize Data and Autonomous Defense to Survive the Next Generation of Cyberattacks

By Nic Adams, Co-Founder & CEO, Orcus

Enterprise security faces a watershed as AI tools mature from passive analytics to autonomous operatives in both offense and defense. To date, traditional firewalls, intrusion detection, signature databases, and manual and/or automated incident response cannot withstand malicious adversaries who brandish self-learning exploit frameworks. Only forward-looking organizations profitably deploying AI at scale, coupled with resilient data architectures and instantaneous behavioral analytics, can achieve determinably lasting security. Key administrators, decision-makers, policy-makers, and powerbrokers must inherently progress beyond experimentation; needing precise methodologies. All of which adequately interlink machine learning, powerful data pipelines, and up-to-the-minute security techniques in order to stay ahead of relentless, persistent threats.

Adversarial AI Agents as Offensive Vanguard

Offense and defense now hinge on an equivalent currency of data. Adversarial AI agents automate reconnaissance, exploit discovery, vulnerability chaining, and binary fuzzing at turbocharged velocity; engineered zero-day exploits by traversing codebases, conducting rapid fuzz campaigns, and infiltrating network protocols concurrently. Above-mentioned agents bank on: self-directed exploration, adaptive reinforcement, emulated user behavior, encrypted payload orchestration, and clandestine command and control (C2) channels to evade conventional detection. Enterprises must presume adversaries train AI models on exfiltrated telemetry, ultimately forecasting defensive logic and custom-tailored attack vectors. As the counter, security teams should field their own AI agents that execute continuous red team simulations, contrive synthetic adversarial scenarios, and chart potential breach trajectories before any malicious actor can punch.

Data Infrastructure as the Nervous System

Effective AI-driven defense brackets on a data infrastructure capable of ingesting, processing, and analyzing vast sensor streams from endpoints, networks, identity systems, and cloud workloads. Data infrastructure managers must construct scalable, low-latency pipelines that incorporate the following: event normalization, feature extraction, real-time indexing, contextual enrichment, and multi-tenant isolation. Hereunder, pipelines feed analytics engines, correlating disparate signals to detect anomalies; outliers in process execution, unexpected data egress events, and/or unusual API call sequences. Without the given foundation, AI models lack fidelity requisite to distinguish benign fluctuations from sophisticated campaigns. Organizations should architect a unified data platform harmonized with time-series logs, behavioral telemetry, and threat intelligence; which becomes the source of truth for both AI model training and on-the-fly intrusion detection.

Securing LLMs and Automated Workflows

Large language models (LLMs) and so-called AI chieftains institute novel attack surfaces. Adversaries exploit prompt injection, model inversion, and training data poisoning to extract proprietary algorithms or manipulate output. Proper risk mitigation entails enterprises implementing resilient guardrails around model inputs and outputs: input sanitization, anomaly detection in prompt patterns, and strict access controls on model endpoints. Additionally, AI orchestrators that schedule and deploy models across hybrid environments must employ attestation mechanisms to verify model integrity at run time; any unauthorized modifications, be it data drift or latent bias insertion, must trigger immediate rollback. By massaging cryptographic signatures into model artifacts and enforcing indispensable distributed validation checkpoints, organizations can prevent silent compromises that otherwise enable downstream breaches and/or data exfiltration.

Actionable Recommendations for CIOs

- **Continuous Adversarial Testing:** Deploy AI-driven red/black teams that simulate multi-vector attacks; API fuzzing, credential stuffing, real-time privilege escalation, and remote code execution. Simulations must run at nanosecond intervals, with rich telemetry generated for model refinement and threat anticipation.
- **Invest in a Unified Data Fabric:** Build a data lake architecture that consolidates on-prem, cloud, and edge data streams. Integrate distributed consensus algorithms for data validation and employ real-time ETL (extract, transform, load) processes to feed AI models with fresh, adversary-resistant data.
- **Seasoned Model Governance:** Institute austere protocols for model lifecycle management; secure model training environments, immutable training data versioning, endpoint access logging, output audit trails, and automated forensic snapshots. The given governance framework prevents undetected model tampering and maintains trust in AI-driven decisions.
- **Resource Allocation for High-Value Projects:** Rather than scattering efforts across dozens of pilots, concentrate on AI initiatives directly enhancing security posture (automated threat hunting, dynamic deception grids, predictive vulnerability scanning, autonomous patch management). Channeling resources into these critical areas aids organizations in achieving faster time-to-value and measurable ROI.
- **Cross-Functional Collaboration:** Split up silos between security teams, data engineers, and AI developers. Establish joint war rooms where threat intelligence, data pipelines, and model performance metrics converge; the environment accelerates decision cycles and reduces response times when new threats surface.

Role of AI in Enterprise Analytics

AI's impact extends beyond security, transfiguring enterprise analytics by sanctioning granular insights into operational patterns, user behavior, and market anomalies. Data infrastructure managers should echelon predictive analytics atop real-time streaming; for example, anomaly detection models, parsing clickstream data, can flag suspicious user sessions. Potentially identifying insider threats or automated bot campaigns before revenue or data integrity suffers. Furthermore, consolidating ensemble learning techniques with graph-based analysis aids organizations in tracing lateral movement paths in networks, linking seemingly unrelated events, and forecasting attacker intent. Habituated insights input back into security controls, creating a closed-loop system where analytics inform defense and defense reinforces analytics.

Amalgamating Security and Data Teams

Security decision-makers must collaborate with data infrastructure managers to ensure that AI-driven defenses do not outpace the underlying data foundation. The common denominator remains, shared metrics: model confidence scores, data freshness indicators, false-positive rates, and incident response latencies, to guide continuous improvements. Regular joint exercises should include stress tests where

data pipelines are intentionally flooded with adversarial examples; practices validate that anomaly detection models maintain high precision without succumbing to signal overload. Embedded security requirements within every stage of data pipeline design boost organizations by preventing critical blind spots and avoid catastrophic lapses in threat coverage.

An Offensive-First Posture Fueled by Data

AI adoption is no longer about vanity projects; offensive-powered security, with a resilient data backbone. Enterprises who master the conjunction of AI agents, next-generation offensive techniques, and data-driven analytics will maintain a sovereign advantage in an everblooming threat landscape. Decision-makers need to eschew complacency and embrace continuous experimentation, steered by adamant data validation and adversarial testing to build a defense platform that preemptively anticipates and neutralizes threats. Reactive controls must become proactive shields that withstand even the most sophisticated assaults.

About the Author

Nic Adams is Co-Founder & CEO of Orcus, the first privatized U.S. commercial hacking startup built by elite black hats, using real-world adversarial experience to outpace nation-state threats and redefine modern cybersecurity. A security architect with black hat hacking roots in non-attributable operations and offensive threat design, he has advised national security stakeholders and private sector leaders on advanced exploitation methodologies and AI-driven attack surfaces. His work focuses on building proactive security systems modeled on real-world adversary capability. Nic can be reached online at X <https://x.com/n1c1337> or LinkedIn <https://linkedin.com/in/nxadams> and at our company website <https://www.Orcus.com/>





Threat Actor Profiling

The art of uncovering adversaries

By Diyar Saadi, Computer Security Researcher, Independent Researcher

Who are the Threat Actors? When it comes to cybersecurity, threat actors aren't always external to the organization; they could be internal or made by the organization itself. There are some cases in which threat actors lie in our actions or systems. Click incursions, for example, by Jerrey's from an email attachment, may lead to ransomware attacks, which could entrap the entire network of an organization. This example accentuates the sophistication surrounding cybersecurity issues in the present. Cyber threat actors are users, groups, or organizations that initiate cyberattacks. These attacks may be rendered to damage, disrupt, illicit, or transact business, and in some cases, cyber terrorism which can be for political, social, or ideological motives. In the past, attacks were executed out of boredom and for fun. Today, cyber attackers are much quieter and the ways in which they strike have advanced immensely.

Types of Threat Actors and Their Characteristics:

Advanced Persistent Threat (APT):

These are groups sponsored by the government, as they are well organized, backed financially, and capable of carrying out long term strategic planning. They can fully resource and fund the operations and exercise a lot of patience, which helps them achieve their goal. Their goal usually is to perform espionage or sabotage.

Financially Motivated Advanced Persistent Threat (FIN):

These actors focus primarily on financial gain. Like other APTs, they too are state-sponsored and well-organized. Their focus on financial systems allows them to exploit weaknesses within entire infrastructures. They are capable of sustained, targeted, and precise attacks over long periods of time.

Hacktivist :

Hacktivists are driven by political or ideological reasons. They are usually less organized than state-sponsored groups. Their goal is to promote a cause or disrupt their targets as a form of protest.

Leaktivists:

Leaktivists are motivated by political or ideological reasons. They specifically aim to obtain and leak sensitive or classified information. Their goal is to reveal corruption, injustice, or secrecy.

Spiders:

These are cybercriminals driven by profit. They often run schemes like Ransomware-as-a-Service (RaaS) or Malware-as-a-Service (MaaS), renting out tools or platforms to other criminals. Sometimes, they are also hired by nation-states for certain operations.

What is Threat Actor Profiling? Threat actor profiling is the process of identifying attackers in cyberspace. It involves understanding the different types of threat actors and their behaviors. However, this profiling is difficult and complex because of the wide range of tools and research from companies and researchers on the methods and techniques used by other groups. Attackers can exploit cloning and copying the techniques of others, making changes or improvements to fit their needs.

Threat actor profiling involves several key categories:

- **Personal Information:** This refers to profiling based on the identity of the attacker, whether an individual or a group. This category provides insight into who is behind the attack and can include factors such as known aliases, past activities, and affiliations. By understanding who the attacker is, defenders can start to anticipate their strategies and behaviors.

- **Intent:** This focuses on understanding the motivation or goal behind the attack. This category examines why the attacker is targeting a particular organization or system. It looks at factors such as financial gain, espionage, hacktivism, or other personal or ideological motivations. Understanding intent helps in predicting future actions and identifying potential targets.

- **Geopolitical Context:** This involves location-based profiling, often used by law enforcement or intelligence agencies. It considers the region or country of origin of the attack and examines whether geopolitical tensions, national interests, or strategic objectives influence the threat actor's behavior.

- **Tools & Methods:** This focuses on the specific techniques, tactics, and procedures (TTPs) used by the attacker during the attack. This category looks at the software tools, malware, and attack vectors used to breach defenses.

- **Target & Assets:** This refers to identifying which systems, networks, or assets the attacker is targeting. This category helps determine what the attacker values and why certain assets or infrastructure are under attack. It provides a clearer picture of the attacker's objectives, whether they aim to disrupt operations, steal data, or cause other forms of damage.

Why Profiling of Threat Actors? Threat actor profiling offers its own benefits, importance, and features. When explained in detail, these can be outlined as follows:

- **Enable Legal Action:** Profiling helps identify and track cybercriminals, leading to arrests and prosecution. By understanding the identity and methods of attackers, law enforcement agencies can take legal action against them and bring perpetrators to justice.

- **Understand Methods:** Profiling provides valuable insights into the tactics, techniques, and procedures (TTPs) used by attackers. This knowledge is essential for developing rules, tools, and products to detect and defend against similar attacks. It helps organizations see the patterns of harmful activity, leading to better preparation and response strategies.

- **Disrupt Attacks:** Profiling can hinder or delay attackers by exposing their methods, which reduces their ability to operate anonymously. By understanding how an attacker works, defenders can implement countermeasures that disrupt the attack process, potentially preventing or reducing the damage caused.

- **Protect Others:** Threat actor profiling allows for information sharing that helps other organizations strengthen their defenses. By sharing knowledge about attack methods, tools, and tactics used by attackers, organizations can better prepare themselves and avoid becoming victims of similar threats. This collaborative approach boosts overall cybersecurity across industries.

Challenges of Profiling Threat Actors? Confusion is one of the most powerful weapons used by adversaries and threat actors in cyberspace. Many other factors can make analysts feel trapped in a maze where everything looks the same and there's no way out. A question arises: Can we track attackers or a specific group by comparing their code to that of another cybercrime group? Unfortunately, the answer is often hypothetical rather than definite. Most cyber attack groups tend to copy techniques and methods from other groups, including coding styles. This makes it harder to attribute attacks with certainty:

Switching Languages: Adversaries may switch languages when communicating across borders. This is commonly seen in forums and dark markets during discussions.

Manipulating Timestamps: Attackers can change timestamps and add false information in the metadata of programs or files. This can be used for a targeted attack.

Mimicking Code: Malicious code for macros or payloads may be designed to look like the behavior of another group's code from a different country.

Changing Time Zones: Attackers may adjust time zones to match the local time of another country. This can confuse analysts about the actual source of the attack.

Using Multiple Proxies: Attackers use several proxies to hide the true source of their attacks. This helps obscure their location.

Domain Registration: A domain may be registered in someone else's name or in a different country. This is done to conceal the attacker's identity.

What are their techniques? Adversaries use a variety of techniques and operational methods that can differ greatly from one group to another. For example, financially motivated groups like FIN tend to target the economic and banking sectors. On the other hand, hacktivist or Leaktivist groups often focus on the public sector, motivated more by political or ideological reasons than by financial gain. This difference also applies to Advanced Persistent Threats (APTs). What sets these groups apart is not just their goals but also their behavior patterns.

Consider cybercriminal factions known as "Spiders." Their main goal is to disrupt or disable systems for financial gain. One of the biggest challenges for cybersecurity researchers is attribution. As discussed in the first part of this series, a group can change its identity, tools, and motives over time. When the players change, even if the techniques and tools stay the same, everything else objectives, motivations, and incentives can shift dramatically. An APT group might start by trying to extract sensitive intelligence but later switch to destructive actions. Similarly, a financially motivated group like the Spiders may suddenly pursue political or ideological aims. In this article, we will examine these patterns with minimal false positives. Later, we will look at why this approach leads to fewer mistakes.

What is the Threat Profiling Pyramid? Until now, the Pyramid of Compromise, also called the Pyramid of Pain, has been a common framework in cybersecurity. But is it still useful in today's world? The Pyramid of Pain has served the cybersecurity community well, but it seems to be outdated.

Traditionally, the pyramid focuses on several key indicators:

1. Hash
2. IP Address
3. Domain Names
4. Network/Host Artifacts
5. Tools
6. TTPs (Tactics, Techniques, and Procedures)

Researchers have relied on this pyramid to gather information about malware. But let's consider some logical questions:

- What if the attacker adds a null byte to the hash?
- What if the attacker uses Fast Flux or Domain Generation Algorithms (DGA)?
- What if the attacker copies another group's TTPs or creates a new technique?

The Pyramid of Threat Actor Profiling is a modern cybersecurity framework designed to offer a behavioral approach to understanding and addressing cyber threats. Unlike traditional models like the Pyramid of Pain, which focuses on **Indicators of Compromise (IOCs) and TTPs** that attackers can easily change, the Pyramid of Threat Actor Profiling goes deeper by examining who the attacker is, why they act, and how they carry out their attacks.


The Pyramid of Threat Profiling Framework:

- **Who:** Identifies the threat actor or group behind an attack.
- **Why:** Understands the threat actor's goal.

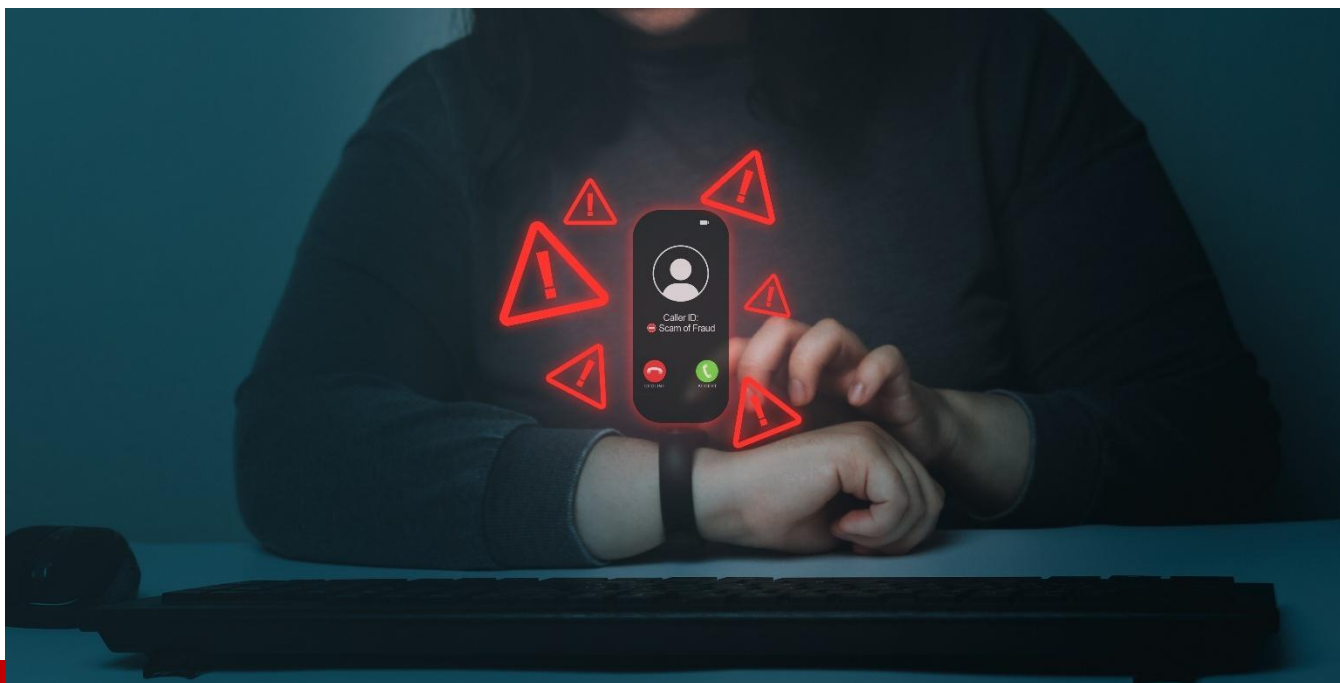
- **What:** Pinpoints the systems and data targeted by the threat actors.
- **How:** Analyzes the techniques, tools, and methods used in the attack.
- **When:** Examines the timing, patterns, and cycles of attacks to predict future threats.

About the Author

Diyar Saadi is a Computer Security Researcher and an independent consultant examining cyber threats. He campaigns against malware while also offering classes on how to reverse engineer it. Diyar has been involved in projects like MITRE ATT&CK Framework and Unprotect.it, concentrating on adversary-focused techniques and tactics, including counter techniques and defense planning. He is also a speaker in many forums around the world, making his contributions to the cybersecurity community. In addition, he has a blog where he discusses reverse engineering and cyber threat intelligence.



Diyar Saadi Ali can be reached online at LinkedIn: <https://www.linkedin.com/in/diyarsaadi/>



Modern Risk, Modern Response: Federal Cybersecurity Needs a Compliance Wake-Up Call

By Peter O'Donoghue, CTO at Tyto Athene, and Gaurav Pal, CEO at stackArmor (a Tyto Athene company)

For years, federal agencies have struggled with the growing tension between innovation and compliance. As cyber threats evolve and emerging technologies like AI become mission-critical, legacy compliance frameworks like [NIST's Risk Management Framework \(RMF\)](#) have increasingly come under scrutiny. While RMF was built to safeguard systems and data, a broad perception is that has become burdened by complexity, slowing down the adoption of game-changing technologies.

Today, the Authority to Operate (ATO), once a badge of cyber-readiness, has become a bottleneck. Government agencies committed to serving and meeting mission operations are frustrated by delays, mounting paperwork, and manual processes that no longer align with how contemporary systems are built, deployed, or maintained in the modern era.

But there's a better way – it doesn't require throwing out the framework, but instead looks at modernizing how agencies implement this critical process.

The Dilemma

The RMF exists to manage risk and not generate more of it. In its current form, the process is often seen as an impediment rather than a mission enabler. Teams across the federal IT landscape have [called for simplification](#), suggesting self-attestation or replacing controls with risk scoring. Others are ready to dismantle the framework altogether.

There is no need to choose between speed and security. Instead, the RMF must evolve to match the pace of mission by embedding automation, infrastructure-as-code, and real-time visibility into every stage of the compliance lifecycle.

Challenges Facing Federal Risk Management

Agencies face growing pressure to deliver secure services with fewer resources. The result? Long-standing compliance challenges become even more difficult to address. The four biggest pain points include:

- Compliance and audit complexity: Documentation-heavy, checklist oriented, manual, and often redundant, the current RMF requires ongoing coordination and exhaustive validation, further burdened by time-consuming audits.
- Limited risk visibility: Disparate systems, inconsistent documentation and the infrequency of assessing controls obscure the true risk picture, leading to decision-makers lacking timely insights into security control effectiveness.
- Velocity: Modern systems evolve continuously driven by DevOps, agile development, and infrastructure-as-code. But compliance often lags, measured in months when systems are changing by the hour. Supporting RMF in this era requires compliance postures to evolve with the systems themselves, including always-on monitoring of key controls in real-time.
- Budget pressures: With agencies under mandates to redirect spending toward mission enabling services, internal support functions like security and compliance often take the hit, threatening long-term resilience and security readiness.

A Smarter Approach to RMF

Agencies don't need to abandon the RMF. Overhauling the RMF approach starts with automation.

Using secure DevOps automation, built with policy and infrastructure as code, government agencies can quickly deploy pre-configured, compliant environments on all major cloud platforms. These environments include embedded security services such as vulnerability management, access controls, and log monitoring mapped directly to [NIST SP 800-53](#) controls.

An approach that leverages a component definition registry further supports this model, offering reusable security components that are validated automatically and ready to plug into multiple systems. When these

tools are deployed “in-boundary,” they reduce external system dependencies and shrink the security footprint.

Even more importantly, this setup enables automated generation of ATO documentation based on real-time system state. That means no more outdated or speculative System Security Plans (SSPs), and no more trying to document security postures that don’t reflect reality.

The Benefits: Faster, Smarter, and More Scalable

The shift to automated compliance unlocks real value for federal teams. Speed to compliance is a massive benefit – the ability to automate the ATO process slashes timelines, cutting months off traditional workflows and getting mission-critical systems into production faster.

While being developed, having auto-generated artifacts for essential systems is important too. This will ensure documentation and evidence are always current and complete. The “audit-ready” approach will also eliminate scrambling before an audit or chasing outdated spreadsheets.

Automation also provides real-time risk awareness, delivering current insight into the effectiveness of security controls. This allows agencies to make faster, better-informed decisions and reduces exposure and improves response.

An automated approach expands past security – it provides enterprise scalability as the same automated approach can be deployed across cloud providers, mission areas, and agency boundaries to maintain consistency. The benefits don’t stop at the cloud either. Automation and DevSecOps practices can extend to on-premise environments, ensuring full-spectrum compliance coverage.

One federal contractor recently automated its FedRAMP Moderate ATO process for a cloud-hosted mission application. By leveraging infrastructure-as-code, policy-as-code, and automated validation, the team reduced its ATO timeline by 50%, improved audit response times, and delivered real-time visibility into control status.

Utilizing an automated-based strategy, the team delivered faster deployments, lower costs, and continuous risk awareness all without compromising on security standards.

RMF Must Evolve with the Mission


NIST’s security controls are not the problem as they have proven their worth over time in safeguarding federal systems. The challenge lies in how those controls are operationalized. By embracing automation, shifting compliance into code, and enabling real-time validation, federal agencies can transform risk management from a burden into a strategic advantage.

To meet the demands of today’s cybersecurity and operational environment, federal agencies must reimagine how RMF is implemented:

- **Dynamic Systems Require Dynamic Compliance:** Legacy concepts like big-bang modernization and static documentation are obsolete. Continuous modernization demands that compliance keeps pace with every system iteration.
- **Automation Is the Glue:** CI/CD pipelines, security as code, and real-time monitoring can eliminate the friction between development and compliance, reducing manual work and increasing consistency.
- **Compliance Must Be Built into DevOps:** Modern development teams can generate machine-readable artifacts using languages like OSCAL, JSON, or YAML that reflect the live state of systems. These artifacts are the foundation for demonstrating continuous, credible compliance.

Collaboration between government and industry is key to successfully modernizing RMF. Technology vendors must not only provide platforms, but they must also bring deep knowledge of federal compliance and a commitment to practical, results-driven approaches. When combined, this will enable federal agencies to move at mission speed, with mission-grade security.

About the Authors



Peter O'Donoghue, CTO, Tyto Athene

As CTO of Tyto Athene, Peter O'Donoghue has over two decades of experience in the technology industry, a history of executive leadership positions at several Fortune 500 companies, and a track record of success in cloud computing, cybersecurity, application dev, and IT modernization.

Recognized by WashingtonExec as one of the Top CTOs to Watch in 2025, his focus at Tyto Athene has been on building a CTO organization aligned with the company's strategic goals, with an emphasis on embedding compelling solutions in new business pursuits and supporting strategic M&A activities. Peter helped lead the acquisition of MindPoint Group, adding top cyber and automation capabilities, rebuilt Tyto Athene's New Business Solutions team, launched an innovation lab, and created a strong technical bench to compete the best and brightest around the Greater D.C. area. Peter is an alum of Leidos, CSC and Unisys.

Peter's eye for building solutions that align to where the technology puck is moving, and his ability to drive the development of new products and services to drive growth, ensures that Tyto Athene remains at the forefront of driving meaningful innovation that achieves mission advantage. To learn more, please visit Peter on LinkedIn at <https://www.linkedin.com/in/peter-o-962154/> and Tyto Athene online at <https://gotyto.com/>.



Gaurav “GP” Pal, CEO, stackArmor (a Tyto Athene company)

Gaurav “GP” Pal is the CEO and founder of stackArmor (a Tyto Athene company). He has over 20 years of digital transformation experience with large U.S. government, defense, healthcare and commercial organizations providing mission acceleration and assurance using commercial cloud services and automating the ATO process. He is a subject matter expert in helping organizations rapidly meet FedRAMP, FISMA RMF, DOD RMF and CMMC compliance. GP received the GCN Rising Star, Fed100 and Meritalk Cyber Defender awards for his role in successful cloud and security program implementations at U.S. federal agencies. He is currently a Board

Member at the Alliance for Digital Innovation (ADI) and Industry Project Leader for the ACT-IAC ato-as-code initiative.

Through GP’s experience and past performance, he is passionate and highly-skilled with a personal commitment to excellence and integrity – consistently building positive and mutually beneficial relationships with partners and government agencies. Connect with GP on LinkedIn <https://www.linkedin.com/in/gppal/> and visit stackArmor online at <https://stackarmor.com/>.



Video Monitoring Security

By Milica D. Djekic

Applying and implementing video surveillance within some property or technological asset is a key pillar of such area's physical safety and security. This system can guarantee better assurance for people, processes and infrastructures, providing truly convenient conditions to all of them. IP cameras are very cheap and still quite reliable equipment that is used for monitoring purposes, offering real-time observing and recording of such a spot. In other words, many incidents can be prevented and if anything happens rigid evidence are obtained via filming, giving a chance to process such a case, as well as deal with something convincing on court. On the other hand, skilled and experienced security professional might assess risks to that infrastructure, simply completing risk assessment sheet. Of particular interest in such a case is to go hand-to-hand with cutting-edge technologies, using them to accelerate data processing and transfer some of human skills to machines. Indeed, nowadays there is tendency to make such transitions relying on AI that is yet in pretty light phase of its development and deployment. Process of training machines to be cognitive as senior security practitioners are is truly time-consuming and dedication-seeking, demanding from R&D teams to recognize, cover and predict all possible scenarios


within their technical project. Human can absorb a lot of that through experience, making those who can provide decision how some level of protection will be estimated and graded to get a pivotal role in such a protocol. Transitioning expert's knowledge from human to machine is extremely difficult task that will not necessarily mean machines could develop thinking process as humans can. AI field is currently in its booming epoch as such solutions once being trained to resolve some complicated problems can do so faster than any human. With growth and progress of emerging technologies everything might appear as much easier, unleashing a plenty of opportunities to communities, economies and industries across the world.

Industrial assets are mainly vulnerable to physical threats, brining on surface a need for good video monitoring security. This means any such facilities should be covered and protected in a physical manner, suggesting significance of surveillance and security practitioners on that place. Capturing all or majority of physical angles and corners with video monitoring setups is a crucial challenge in assuring such a space. This matters in sense of prevention, online monitoring and incident response – very similarly as in case of virtual safety and security. Today many properties, assets and large infrastructures are protected using IP cameras that are quite inexpensive and dependent on web signal connectivity, making them also sensitive to cyber-attacks. Accordingly, these networks also have a prefix “cyber” and their functionality and operability strongly require stable internet connection, truly being critical infrastructure in ongoing connotation. Security staff working on such a set must take into consideration even those conditions, carefully completing risk assessment sheet and determining all strengths and weaknesses of that complex. It's very hard developing criteria of thinking, judging and marking all risky elements, putting all those requirements into helpfully created spreadsheets that must be examined in truly rational fashion. In order to deal with such risk assessment sheet security practitioners are supposed to pass through vetted training and even that is not sufficient as they also need to deserve some experience, making their skill better trusted. Senior security practitioner knowledge is literally priceless, opening up many outlets to machine intelligence designers to make next-generation systems that can handle something complicated in time-effective manner.

For instance, some of AI-powered solutions can mimic or even demonstrate skill humans have, but they cannot cope with their own thoughts and other cognitive behavior capabilities as they are not yet at such scale of their development. Practical experience with AI indicates it's needed to imagine, explain and make good problem-solving of some environment, safeguarding great orienteering of all elements of that surrounding. In other words, if navigating in such a space is maximized that technical solution will deal with higher IQ to such sets of tasks, streamlining that system to go beyond intelligent or even smart.

About The Author

Milica D. Djekic is an Independent Researcher from Subotica, the Republic of Serbia. She received her engineering background from the Faculty of Mechanical Engineering, University of Belgrade. She writes for some domestic and overseas press and she is also the author of the books *"The Internet of Things: Concept, Applications and Security"* and *"The Insider's Threats: Operational, Tactical and Strategic Perspective"* being published in 2017 and 2021 respectively with the Lambert Academic Publishing. Milica is also a speaker with the BrightTALK expert's channel. She is the member of an ASIS International since 2017 and contributor to the Australian Cyber Security Magazine since 2018. Milica's research efforts are recognized with Computer Emergency Response Team for the European Union (CERT-EU), Censys Press, BU-CERT UK and EASA European Centre for Cybersecurity in Aviation (ECCSA). Her fields of interests are cyber defense, technology and business. Milica is a person with disability.





When Theory Meets Reality: The Fatal Flaws in Traditional Incident Response

By Andy Lunsford, CEO and Co-Founder, BreachRx

Incident response is broken.

For years, organizations have been investing: tools, training, tabletop exercises, and everything in between. You name it, security teams have tried it. But at the end of the day, most organizations are not prepared to respond effectively when a cyber incident hits. Plans may look solid on paper. CISOs will tout comprehensive playbooks and documented policies ready to be executed by cross-functional teams. However, plans often fall apart when theory meets reality.

Good intentions often drive traditional approaches to incident response (IR). But they are plagued by execution gaps. The idea of a playbook is excellent, but if it's treated as a loose guideline and not an actionable, auditable checklist, it won't be effective. Teams too often rely on tribal knowledge or gut instinct instead of implementing repeatable processes. And when the rubber meets the road, teams

consistently struggle with fragmented communications and limited visibility -- issues that inflate incident timelines.

The current IR crisis starts with the way plans are designed and executed. You'd be shocked at how many static documents or spreadsheets I've seen companies try to pass off as an IR plan. These are documents that don't integrate into actual workflows. These files will sit on a server and will help a company pass audits, but that's about the extent of their utility. They're not doing much to support response teams in the moment. In real-world incidents, teams need more than just a set of instructions.

They need a dynamic, shared environment where work is tracked in real-time, responsibilities are clearly defined, and progress is visible across the organization. Without that as a baseline, efforts will inevitably become siloed and inconsistent at best (and counterproductive at worst). When the left hand doesn't know what the right hand is doing, critical tasks will fall through the cracks.

Plans are Not Preparedness

Another issue is the false sense of confidence that traditional IR plans can create. Organizations conduct tabletop exercises that are, essentially, storytelling sessions. Participants gather in a room (or on Zoom), discuss a hypothetical scenario, and check a box indicating that the plan was tested. However, real-world incidents don't unfold in two hours. They evolve over days or weeks. They require coordination across security, legal, IT, HR, line-of-business executives, and communications teams.

They demand live decision-making, systems access, and continuous updates – and that surely doesn't fit neatly into a pre-scheduled two-hour window. They can (and will) take place at the worst possible time: when your CISO is sleeping, when your General Counsel is on vacation, or during your busiest time of year. A tabletop won't tell you whether half your team lacks access to key tooling, or whether the legal team and the SOC are aligned on notification timelines. You only discover those gaps when you practice with real systems, fundamental roles, and real timelines.

Modern incident response requires a shift in mindset. Static documentation must transform into active coordination. Effective response requires a dynamic command and control approach. You have your primary objective (neutralizing and eliminating the threat). Still, dozens of side quests must be managed along the way: security analysts investigating alerts, IT staff remediating systems, legal and privacy teams handling notification requirements, executives assessing the business impact, and employees on the front lines managing the company's reputation with customers. The list goes on.

These side quests branch out unexpectedly, they loop back, and they often require different teams to act independently while staying aligned. For that to work, there needs to be a common operating picture: one place where tasks are assigned, updates are logged, decisions are documented, and everyone involved can see the full scope of what's happening.

Yes, this presents technology challenges. But at its core, this tends to be a human challenge. In my experience, the most overlooked factors in IR are trust, relationships, and communication. Incident response teams aren't just lists of names in a plan; they're people who need to work together under extreme stress. The best way to build that capability is to invest in relationships before a crisis occurs.

Grab coffee with your legal counterpart. Walk through a scenario with comms and HR. Take the time to understand what your executive team worries about most when it comes to cyber risk. These touchpoints build trust, and trust is the currency of effective collaboration when the stakes are high.

From Chaos to Coordinated Execution

When trust is lacking, or when teams haven't practiced working together, even the best-laid plans can unravel. One company we work with learned this the hard way. After suffering a ransomware attack, they followed their documented process as closely as possible. However, the outage persisted for days, and they struggled with coordination, reporting, and decision-making.

Later, when evaluating new approaches, they simulated that same incident using a modern response framework. The difference was stark. They didn't just talk through what they would do: they did it. They assigned real tasks to real people, documented actions, and practiced with the actual tools they'd rely on in a live event. Their general counsel called it the best exercise they'd ever run, not because it was smooth, but because it wasn't. It surfaced real issues they could now fix.

This company learned a very basic lesson: that resilience isn't built through theory. Resilience comes from experiences with realistic, scenario-driven practice that tests your limits. A resilient organization is one with frameworks that are informed by—and can adapt to—the twists and turns of real-life incidents. A resilient organization is built through a culture where IR is not a once-a-year update to that doc on the server: it's an ongoing, evolving capability that spans tools and teams.


Leadership, Not Lone Wolf

To be clear, this doesn't mean CISOs and security leaders need to have all the answers. One of the most common failure modes we've seen is the leader who feels they must solve everything themselves. That's not sustainable, and it's not a strategic approach. The role of the CISO during an incident isn't to be a hero. It's to be the coordinator, the communicator, and the bridge between technical response and business decision-making. That means knowing where to turn for expertise – whether it's internal IR analysts, outside counsel, or third-party support – and having those relationships in place before the fire starts.

The unfortunate reality is that most organizations will face a cyber incident at some point. The goal of eliminating every risk is impossible. Instead, the goal should be to put your organization in a position to respond more effectively, faster, and with greater confidence. Build guardrails against chaos. That means frameworks that support real-time coordination, the ability to adapt to dynamic situations, and a modern approach to IR that reflects how organizations work today.

About the Author

Andy Lunsford is the CEO and Co-Founder of BreachRx, the provider of the first intelligent incident response platform designed for the entire enterprise. Prior to founding BreachRx, Andy spent 15 years in privacy law and large-scale commercial litigation. Andy co-founded BreachRx to transform incident response and reporting into a routine operational business process while shielding C-level executives from personal liability. Andy has a BA from Washington and Lee University, a JD from the University of Arkansas, and an MBA from the Wharton School of the University of Pennsylvania.



Andy can be reached online at alunsford@breachrx.com and at our company website <https://www.breachrx.com>.



Where Do Leadership and Technical Teams Disagree on AI-Driven Network Security?

By Mitch Densley, Principal Solutions Architect at Opendgear

Digital transformation brings new opportunities, but it also introduces complexity. Today's Artificial Intelligence (AI)-powered systems take in and produce data at a scale that would have been hard to imagine just a few years ago. Digital transformation also introduces new vulnerabilities that jeopardize network integrity, granting attackers access to subsets or supersets of data. These challenges demand solutions and technologies that ensure network resilience, security and efficiency. While not a cybersecurity silver bullet, AI helps humans spot needles in haystacks and enables them to prioritize risk and incident response activities.

Although there is almost universal agreement that AI will play a crucial role in network security in the future, leadership and technical teams have different perspectives about the nature of its implementation, among other topics. A research report that surveyed 513 Chief Information Officers (CIOs) and Chief Security Officers (CSOs), including 508 network engineers across Western countries, uncovered key discrepancies between both groups concerning AI. Organizations must reconcile this misalignment to fully realize AI's potential in securing their networks.

Differing Views on AI Implementation

A closer look into the research reveals that the levels of reported AI adoption vary considerably between leadership teams and network engineers. CIOs and CSOs had an optimistic view of AI adoption, with a majority citing full implementation. However, network engineers reported a lower level of integration, with only 42% claiming their businesses had fully implemented AI. Another important finding concerned the level of AI integration into existing cybersecurity infrastructure. Only 28% of surveyed engineers said AI had been integrated fully, while 63% reported integration to some degree but not fully.

One reason for this difference could be that leadership focuses on broader strategic goals, whereas network engineers assess AI based on its direct impact on network performance and security. Engineers view the world through a granular lens, while leadership focuses on broader strategic goals and generalities. Put another way, if business leaders were asked, "What's on McDonald's menu?" They'd say "hamburgers" or "burgers and fries." Engineers would probably look up the menu and respond with, "burgers, fries, chicken, salads, desserts, etc."

Another interesting point of contention had to do with who should lead AI rollouts in network management. Some proposed roles included the CIO, CSO, head of IT security, or even an AI/ machine learning specialist. This discrepancy further highlighted the diverse priorities and areas of expertise within businesses.

Organizations should establish shared AI deployment frameworks to align clear goals, standard metrics, and mutual input across teams to bridge these gaps. Encouraging consensus in the early stage helps align expectations, leading to a more cohesive implementation strategy.

Resource Allocation and AI Integration Challenges

There are a variety of network management and security applications possible through AI. For example, organizations can use AI to enhance anomaly and threat detection, specifically malware and network anomalies, ultimately preventing breaches. Nevertheless, differing priorities emerged between leadership and technical teams when considering resource allocation for AI-driven network management. The report found that nearly 70% of engineers believe AI will enhance their organization's ability to respond to cybersecurity incidents. However, 66% of CIOs and CSOs allocated only 4-10% of their IT and cybersecurity budget to AI for network management.

In addition to these disagreements between leadership and technical teams, organizations face several hurdles to AI implementation. The most commonly cited challenge from network engineers was the high initial investment required to deploy AI technologies, with 29% of respondents identifying this as the top barrier to AI adoption. This sentiment from engineers mirrors their belief that there aren't enough resources getting allocated to support AI implementation. A close second challenge at 28% was regulatory compliance, followed by the need for skilled professionals to manage AI-driven systems. Concerning the latter, 31% of network engineers said they plan to prioritize training and upskilling efforts with IT staff to manage and integrate AI technologies effectively.

For long-term success, leadership should invest in effective upskilling programs for both engineers and executives. Prioritizing AI training ensures that teams can advance their capabilities while reducing their reliance on temporary fixes, setting the foundation for sustainable and secure AI adoption.

How Out of Band Management Can Support AI-Integrated Networks

An equally intriguing insight from the report was that of those CIOs or CSOs whose organizations are either planning to or have partially implemented AI for network management, 32% (the highest proportion) are preparing for their AI implementation by employing continuous monitoring and real-time analytics. Highly complementary to this context are Out of Band (OOB) solutions, which offer independent and secure access to network assets, enabling engineers to conduct provisioning, orchestration and remediation from remote locations. By creating an independent management plane separate from the production network, OOB solutions permit AI tools to remain operational even if the primary network goes down, simultaneously ensuring real-time analytics use current data.

Organizations investing in AI also need to safeguard that investment with robust OOB solutions. As these systems protect critical infrastructure, data networks, storage, GPU clusters during outages or disruptions, preserving the integrity and uptime of AI operations.

Fostering Alignment Between Leadership and Technical Teams

Despite different perspectives and opinions, both leadership and network engineers agree on AI's potential to improve network management and cybersecurity. Building on this shared recognition, organizations can foster collaboration by aligning objectives, defining shared goals and prioritizing initiatives such as IT training, network resilience and real-time analytics.

Many organizations opt for a step-by-step approach to AI adoption, starting with partial deployment and scaling up as they gain experience and familiarity with the technology. Such an approach allows businesses to test AI capabilities on a smaller scale, providing a proof of concept before a full implementation. By working together to integrate AI effectively, leadership and technical teams can maximize its value in network security and management.

All in all, businesses should combine shared frameworks, funding for workforce development, and infrastructure like OOB to enable scalable, resilient AI. By integrating these specific actions, leadership and technical teams can work together more effectively to realize AI's full potential in network security and management.

About the Author

Mitch Densley is a Principal Solutions Architect at Opendgear, where he designs and implements Smart OutofBand™ management networks that underpin secure, resilient infrastructure. With deep expertise in cybersecurity, network security, and technical training, Mitch applies Zero Trust design and defenseindepth strategies to help organizations stay ahead of disruption. Known for distilling complex concepts into practical solutions, he champions customer success through collaboration with Opendgear's engineering and support teams. His contributions include thought leadership panels and podcast appearances on securing AI and GenAI infrastructure, drawing industry attention to the unique demands of modern computing environments.





EVENTS



HOSTED BY

مجلس الأمن السيبراني
CYBER SECURITY COUNCIL



IN PARTNERSHIP WITH



UAE (Abu Dhabi)
Section

6th MIDDLE EAST INSTRUMENTATION, CYBERSECURITY & AUTOMATION EXPO 2025

OIL & GAS


UTILITIES

HEAVY INDUSTRIES

16-18

September, 2025

Rixos Marina, Abu Dhabi, UAE



KEYNOTE SPEAKER

H.E. DR. MOHAMED AL KUWAITI

HEAD OF CYBER SECURITY,
UNITED ARAB EMIRATES GOVERNMENT

UAE Cyber Security Council (CSC)



www.meicaexpo.com

GRF SUMMIT ON SECURITY & THIRD-PARTY RISK

**NOVEMBER 3-5, 2025
THE PALMS LAS VEGAS**

Networking and Education on Critical Third-Party and Cybersecurity Issues, for Mutual Resilience

The conference features dozens of speakers on third-party risk management, cloud security, emerging cybersecurity threats, and AI/machine learning threat mitigation and management.

Attendees will gain an understanding of how some of the largest and most sophisticated organizations in the world are managing risk, and leave the conference better armed to defend their company, regardless of its size or the status of its risk mitigation program.



Tri-Service Asian Defense & Security Exhibition,
Conference and Networking Event

10-13 November
2025

IMPACT , Muang Thong Thani, **THAILAND**



Power of Partnership

www.asiandefense.com

Shifting Power in Cyber Defense

3rd-5th December 2025 | Hotel Intercontinental, KL, Malaysia

AVAR's annual conferences have been hosted in China, Dubai, Hong Kong, Japan, Malaysia, Singapore, Sydney, and in many other prominent locations in the Asia Pacific region in every year since 1998. AVAR's events are known for exemplary knowledge sharing, arrangements, and hospitality, and are regarded as Asia's most impactful cyber security conferences for CXOs.



Delegates



Speakers



Presentations

AVAR 2025 recognizes the growing importance of interaction between cyber security practitioners and management and includes CISO oriented knowledge sessions and awards.



Panel Discussions



CISO Connect



CISO Awards

Who Should Attend

CEOs	CTOs	CSOs/CISOs	Regulators
Law Enforcement Agencies		Researchers/Professors/Students	

Sponsors

Silver Sponsor



Bronze Sponsor



T Shirt Sponsor



Delegate Bag Sponsor



Attendee Name Badge & Lanyard Sponsor



For more information on AVAR 2025
and Sponsorship benefits, contact

✉ rgdwivedy@aavar.org
☎ +91 93840 19113

www.aavar.org/cybersecurity-conference



CYBER DEFENSE TV

INFOSEC KNOWLEDGE IS POWER

CyberDefense.TV now has 200 hotseat interviews and growing...

Market leaders, innovators, CEO hot seat interviews and much more.

A division of Cyber Defense Media Group and sister to Cyber Defense Magazine.

The Interviews

These anticipated **"CEO Hotseat"** Interviews will feature a C-level executive from the hottest Infosec companies being interviewed by **Gary Miliefsky**. Gary is an internationally-recognized speaker and Infosec expert and will make the interviews lively, informative, and highly favorable to the interviewees.

CYBER DEFENSE TV | © 2018 CYBER DEFENSE MAGAZINE. All Rights Reserved. www.cyberdefense.tv

Free Monthly Cyber Defense eMagazine Via Email

Enjoy our monthly electronic editions of our Magazines for FREE.

This magazine is by and for ethical information security professionals with a twist on innovative consumer products and privacy issues on top of best practices for IT security and Regulatory Compliance. Our mission is to share cutting edge knowledge, real world stories and independent lab reviews on the best ideas, products and services in the information technology industry. Our monthly Cyber Defense e-Magazines will also keep you up to speed on what's happening in the cyber-crime and cyber warfare arena plus we'll inform you as next generation and innovative technology vendors have news worthy of sharing with you – so enjoy. You get all of this for FREE, always, for our electronic editions. [Click here](#) to sign up today and within moments, you'll receive your first email from us with an archive of our newsletters along with this month's newsletter.

[By signing up, you'll always be in the loop with CDM.](#)

Copyright (C) 2025, Cyber Defense Magazine, a division of CYBER DEFENSE MEDIA GROUP (STEVEN G. SAMUELS LLC. d/b/a) 276 Fifth Avenue, Suite 704, New York, NY 10001, Toll Free (USA): 1-833-844-9468 d/b/a CyberDefenseAwards.com, CyberDefenseConferences.com, CyberDefenseMagazine.com, CyberDefenseNewswire.com, CyberDefenseProfessionals.com, CyberDefenseRadio.com, and CyberDefenseTV.com, is a Limited Liability Corporation (LLC) originally incorporated in the United States of America. Our Tax ID (EIN) is: 45-4188465, Cyber Defense Magazine® is a registered trademark of Cyber Defense Media Group. EIN: 454-18-8465, DUNS# 078358935. All rights reserved worldwide.
marketing@cyberdefensemagazine.com

All rights reserved worldwide. Copyright © 2025, Cyber Defense Magazine. All rights reserved. No part of this newsletter may be used or reproduced by any means, graphic, electronic, or mechanical, including photocopying, recording, taping or by any information storage retrieval system without the written permission of the publisher except in the case of brief quotations embodied in critical articles and reviews. Because of the dynamic nature of the Internet, any Web addresses or links contained in this newsletter may have changed since publication and may no longer be valid. The views expressed in this work are solely those of the author and do not necessarily reflect the views of the publisher, and the publisher hereby disclaims any responsibility for them. Send us great content and we'll post it in the magazine for free, subject to editorial approval and layout. Email us at marketing@cyberdefensemagazine.com

Cyber Defense Magazine

276 Fifth Avenue, Suite 704, New York, NY 1000

EIN: 454-18-8465, DUNS# 078358935.

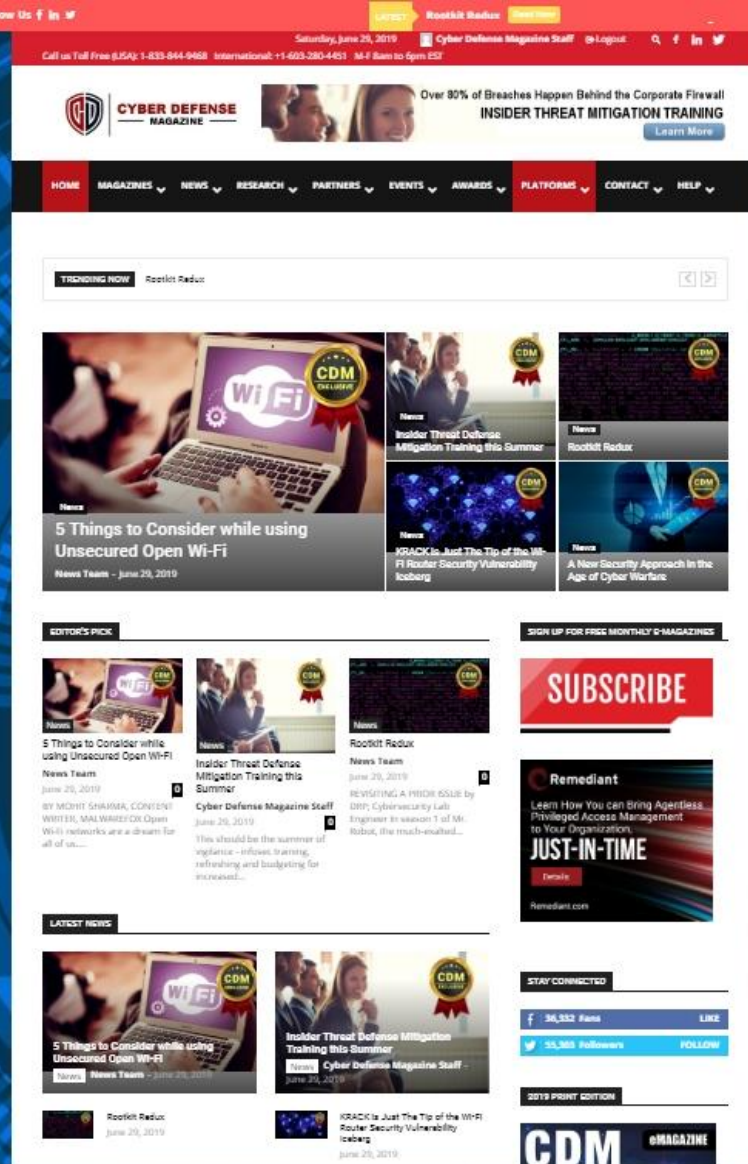
All rights reserved worldwide.

marketing@cyberdefensemagazine.com

<https://www.cyberdefensemagazine.com/>

NEW YORK (US HQ), LONDON (UK/EU), HONG KONG (ASIA)

Cyber Defense Magazine - Cyber Defense eMagazine rev. date: 09/02/2025



Books by our Publisher: [Amazon.com: CRYPTOCONOMY®, 2nd Edition: Bitcoins, Blockchains & Bad Guys eBook : Miliefsky, Gary: Kindle Store, Kindle Store, Cybersecurity Simplified, The AI Singularity: When Machines Dream of Dominion with others coming soon...](https://www.amazon.com/dp/B075N1Y1Y1)

13 Years in The Making...

Thank You to our Loyal Subscribers!

We've Completely Rebuilt CyberDefenseMagazine.com - Please Let Us Know What You Think. It's mobile and tablet friendly and superfast. We hope you like it. In addition, we're past the five nines of 7x24x365 uptime as we continue to scale with improved Web App Firewalls, Content Deliver Networks (CDNs) around the Globe, Faster and More Secure DNS and CyberDefenseMagazine.com up and running as an array of live mirror sites. We successfully launched <https://cyberdefenseconferences.com/> and our new platform <https://cyberdefensewire.com/>

CDM

CYBER DEFENSE MAGAZINE

THE PREMIER SOURCE FOR IT SECURITY INFORMATION

eMAGAZINE

www.cyberdefenseemagazine.com

"Cyber Defense Magazine is free online every month. I guarantee you will learn something new you can use to help you improve your InfoSec skills."

Gary S. Miliefsky, Publisher & Cybersecurity Expert



**ALWAYS FREE
NO STRINGS ATTACHED**



CYBER DEFENSE MAGAZINE

WHERE INFOSEC KNOWLEDGE IS POWER



www.cyberdefensewire.com

www.cyberdefensetv.com

www.cyberdefenseradio.com

www.cyberdefenseawards.com

www.cyberdefenseconferences.com

www.cyberdefensemagazine.com

Shadow IT May Leave You at Risk.

Unauthorized or unknown internet-facing assets
in your network can expose sensitive defense
information to our adversaries.

NSA's no-cost cybersecurity services can
help you find and protect your assets to
better secure your network.

GET STARTED TODAY AT

nsa.gov/ccc





*** with help from writers
and friends all over the Globe.**