LEADERSHIP | STRATEGY | RESILIENCE

The AI-ERA CISO

LEADERSHIP STRATEGY AND PRACTICAL GUIDANCE FOR A MORE SECURE AI-POWERED WORLD

WILLIAM D. SULLIVAN

The AI-Era CISO

A Strategic Guide to Security Leadership in the Age of Artificial Intelligence

Steve Publications

This book is available at https://leanpub.com/theai-eraciso

This version was published on 2026-09-09

Contents

A Strategic Guide to Security Leadership in the Age of Artificial Intelligence

In an era where artificial intelligence amplifies both attack and defense, where autonomous agents outnumber humans in enterprise systems, and where a single compromised software dependency can cascade into a global crisis, the Chief Information Security Officer faces the most consequential transformation in the profession's history. This book examines what it truly means to lead security in this new environment: how the CISO role is evolving from a technical discipline into an enterprise-wide leadership mandate centered on digital trust, organizational resilience, and AI-era security governance. Drawing on frameworks, case studies, regulatory developments, and real-world incidents from the past decade, this guide provides CISOs, boards, and security leaders with the practical knowledge, strategic frameworks, and execution models needed to protect organizations while enabling the innovation they depend on.

Introduction: The CISO at the Crossroads

In January 2024, a finance employee at Arup, a global engineering firm, joined what appeared to be a routine video conference with the company's chief financial officer and other senior leaders. They discussed an urgent transfer. The employee followed protocol, completed the authorization, and wired 25.5 million dollars to accounts controlled by fraudsters. Later that day, the truth emerged. None of the executives on the screen were real. Every face, every voice, every gesture was generated by artificial intelligence. The systems had not been breached, the networks had held, but trust itself had been compromised. [1]

The Arup incident captures something essential about the new era we are entering. It was not a vulnerability in code that caused the loss, not a misconfigured server, not a leaked password. It was a collision between human judgment and machine-generated deception that any moderately funded criminal group could now produce. And it was preventable only through governance, process, and organizational awareness that many CISOs are only beginning to grapple with.

Two years earlier, a different kind of supply chain attack sent shockwaves through the open source community. A Microsoft engineer debugging performance anomalies in PostgreSQL noticed something odd in SSH logins: extra CPU cycles, suspicious delays. What he uncovered was a backdoor planted in XZ Utils, a compression library installed on virtually every Linux system in the world. The attacker had spent years gaining trust in the project, methodically introducing changes that appeared legitimate while slowly building a foothold capable of compromising millions of servers. The fact that it was discovered at all depended on luck, not architecture. [2]

And in February 2024, the Change Healthcare ransomware attack demonstrated how a single vendor compromise could destabilize an entire industry. The Russian-affiliated BlackCat group gained access to Change Healthcare through credentials on a server that lacked multi-factor authentication. From there, they encrypted critical systems that handled roughly half of all U.S.

pharmacy claims and millions of patient payments. The American Hospital Association called it the most significant and consequential incident of its kind against the U.S. healthcare system. UnitedHealth Group would advance nearly 4.7 billion dollars to providers affected by the disruption. [3]

Three incidents, three different vectors, one common reality: the modern threat environment has grown too complex, too fast, and too interdependent for conventional security leadership to manage alone. The CISO who succeeds in this environment must think less like a technician and more like an enterprise risk leader, a strategist, and an organizational architect.

The Moment: Why Now Is Different from Every Previous Era

Security has always been a discipline of managing trade-offs. Every organization has more risks than resources to address them, so leaders must decide what to protect first, what to tolerate, and what to transfer. The basic calculus has not changed. What has changed is the speed, scale, and asymmetry of the threats organizations face.

On the offensive side, artificial intelligence has democratized sophisticated attack capabilities. Generative AI tools now enable threat actors to craft convincing phishing campaigns in multiple languages, generate polymorphic malware that evades signature-based detection, and conduct social engineering at scale with personalization that would previously require a dedicated operations team. A Darktrace survey found that 74 percent of cybersecurity professionals agree AI-powered threats now pose a significant challenge to their organizations. [4] The UK's National Cyber Security Centre warned that AI is already being used by every type of threat actor and will almost certainly increase the volume and impact of cyberattacks in the years ahead. [5]

On the defensive side, AI promises real benefits but introduces its own failures. According to IBM's 2024 Cost of a Data Breach Report, organizations using AI and automation in their security operations saved an average of 2.2 million dollars per incident and shortened detection and containment timelines by 98 days. [6] Yet the same technology also creates new risks: AI platforms ingesting sensitive data without adequate controls, models vulnerable to prompt injection and data poisoning, and autonomous agents operating across enterprise systems with inadequate visibility or guardrails.

Then there is the regulatory and geopolitical environment. The EU AI Act imposes strict requirements on high-risk AI systems and general-purpose models with systemic risk. [7] The SEC requires public companies to disclose material cybersecurity incidents within four business days. [8] New U.S. executive orders address software supply chain security and mandate post-quantum cryptography migration. Across the globe, personal liability for executives overseeing cyber failures is emerging in law and practice, with directors increasingly held accountable for governance failures even when no breach has occurred. [9]

The economics are shifting as well. The average cost of a data breach reached 4.88 million dollars globally in 2024, up 10 percent from the previous year. [10] Cyber insurance markets, which hardened dramatically between 2020 and 2022, have softened somewhat but now demand evidence of robust security controls, including multi-factor authentication, endpoint detection and response, and privileged access management, before underwriting policies. [11] Cybercrime costs are projected to exceed 10 trillion dollars annually within a few years. [12]

Finally, the human dimension has become critical. ISC2's 2025 Cybersecurity Workforce Study, based on 16,029 global respondents, found that 36 percent of organizations report budget cuts, 39 percent face hiring freezes, and 24 percent have enacted layoffs in their security teams. Yet AI expertise is in the top tier of skills shortages, cited by 41 percent of respondents. [13] Meanwhile, CISO burnout is reaching crisis levels. Proofpoint's 2025 Voice of the CISO Report found that 63 percent of CISOs experienced or witnessed burnout in the past year, and only 64 percent feel aligned with board priorities, down from 84 percent the previous year. [14]

All of these forces are converging. The CISO who entered the role a decade ago, focused on firewalls, antivirus, and perimeter defense, would find that role nearly unrecognizable. The modern mandate requires fluency across technology, business strategy, regulatory compliance, organizational psychology, risk economics, and increasingly, artificial intelligence itself.

From Gatekeeper to Growth Enabler

The traditional image of the CISO is that of a gatekeeper: the person who says no when a new product wants to launch without security review, who blocks

access for convenience, who turns risk into a compliance exercise. That model worked in a simpler world where security was primarily about keeping bad actors out of a defined perimeter. But the perimeter no longer exists, and organizations that treat security as a brake rather than an enabler will find themselves unable to compete.

Security leaders are being elevated. Deloitte's Future of Cyber Survey, polling 1,200 decision-makers across 43 countries, found that 73 percent saw increased CISO involvement in strategic technology conversations over the prior twelve months. [15] CISO salaries have risen dramatically, with average compensation exceeding 550,000 dollars and top roles surpassing 1 million dollars annually. [16] Many are being re-titled as Chief Security Officers with broader responsibilities encompassing physical security, business continuity, and crisis management.

Yet this elevation brings pressure, not just prestige. A Trellix survey found that 49 percent of CISOs do not see a future in their roles due to workload and regulatory pressures. [17] Gartner estimates that nearly half of cybersecurity leaders will change roles by the end of the decade, with a quarter transitioning to entirely different positions because of stress. [18] The role demands more from individuals than most organizations are prepared to provide in support, resources, or realistic expectations.

The successful modern CISO has shifted their identity from defender to enabler. They frame security not as a constraint on business but as a prerequisite for business. They embed security into product development workflows rather than reviewing them after the fact. They treat incidents as inevitable and focus on resilience rather than prevention alone. And they speak the language of the board: financial impact, strategic risk, competitive advantage, and regulatory exposure.

This transformation is not cosmetic. It requires a fundamental recalibration of how the CISO thinks about their role, how they structure their organization, how they allocate budgets, and how they engage with stakeholders across the enterprise.

The Thesis: What This Book Will Argue and Deliver

This book is built on a central thesis: the CISO role in the AI era must evolve into enterprise risk leadership centered on three core capabilities.

First, strategic business alignment. The CISO must understand the business the way the CEO does: what drives revenue, what creates competitive advantage, what would constitute a material disruption. Security strategy must emerge from business strategy, not from a catalog of threats. The CISO must be able to articulate, in business terms, why a particular risk matters, what would happen if it materialized, and how much it costs to reduce it.

Second, AI-competent technical depth. You do not need to be a machine learning engineer to be an effective CISO, but you must understand AI systems well enough to govern them. That includes understanding the security risks of AI models, training data, and deployment pipelines. It includes understanding how threat actors use AI, how deepfakes and automated social engineering work, and how AI agents expand the attack surface. And it includes knowing how to leverage AI defensively while recognizing its limitations and failure modes.

Third, organizational leadership across stakeholders, regulators, and ecosystems. Security is no longer something the CISO team does. It is something the organization must do, with the CISO providing vision, standards, and accountability. The modern CISO must influence without authority, build coalitions across functions, manage relationships with regulators and law enforcement, coordinate with third parties and industry groups, and translate technical uncertainty into clear recommendations for decision makers who will bear the consequences of those decisions.

Beyond these three pillars, this book covers the full spectrum of CISO responsibilities in the AI era, including:

- Strategic planning: how to develop, document, and execute a cybersecurity strategy aligned to business objectives.
- Governance: how to structure board oversight, define decision rights, and establish accountability.
- Risk management: how to quantify cyber risk, integrate it into enterprise risk management, and make investment decisions under uncertainty.
- Organization design: how to structure security teams for scale, agility, and effectiveness.
- Talent and culture: how to attract, develop, and retain security professionals while managing burnout and building security awareness across the enterprise.

- Budgeting: how to allocate resources across prevention, detection, response, and resilience, and how to build business cases that executives will fund.
- Core technical domains: identity, cloud, data, applications, software supply chain, security operations, and incident response.
- AI-specific responsibilities: securing AI systems, managing AI adoption risks, defending against AI-enabled attacks.
- Third-party risk: managing vendor exposure, concentration risk, and ecosystem dependencies.
- Communication: how to report to the board, influence executives, and manage crises.
- Resilience: ensuring the organization can continue operating under attack and recover quickly from disruption.
- The future: how emerging technologies, regulations, and geopolitical trends will reshape the role over the next five to ten years.

This book is written for CISOs and aspiring CISOs, but it is also relevant for the stakeholders they must work with: board members overseeing cyber risk, CEOs deciding how security fits into their leadership team, CIOs and CTOs building the architecture the CISO must secure, risk officers integrating cyber into enterprise risk, and legal and privacy leaders navigating the regulatory environment.

How to Use This Book

The chapters are arranged in a logical progression from foundations to advanced topics, though you do not need to read them in order. If you are a current CISO looking to strengthen a specific area, jump to the relevant chapters. If you are building your first comprehensive strategy, read chapters four through nine before diving into technical domains. If you are a board member or executive seeking to understand what a strong CISO mandate looks like, chapters five, six, eighteen, and nineteen will be most relevant.

Each chapter is self-contained but contributes to the whole. You will find frameworks you can adapt, decision models you can apply, case studies that illustrate both failures and successes, and benchmarks you can use to evaluate your own program. The book draws on frameworks from NIST, OWASP, FAIR,

MITRE, NACD, and other recognized standards bodies, but frames them in practical terms rather than compliance checklists.

Some chapters will challenge conventional assumptions. For example, the book argues that zero trust, while valuable in principle, is often over-invested in at the expense of more fundamental controls. It argues that AI adoption in security operations is overhyped in some areas and underutilized in others. It questions whether the industry's obsession with prevention has led to insufficient investment in resilience. These positions are not dogma, but arguments built on evidence, trade-offs, and the lessons of real incidents.

The book does not offer a one-size-fits-all blueprint. A CISO at a 50-person startup faces different constraints and priorities than a CISO at a multinational bank or a critical infrastructure operator. Each chapter includes guidance on how to scale approaches to different organizational contexts.

What you will not find in this book are exercises, quizzes, or discussion questions. This is not a training manual. It is a research-driven, practitioner-oriented treatment of what it means to lead security in the age of artificial intelligence, written to be read, referenced, and applied by professionals who operate at the highest levels of responsibility.

The mandate has changed. The CISO who embraces the transformation will lead an organization that is more resilient, more trusted, and more capable of competing in the digital economy. The CISO who does not will find themselves managing a crisis they were never prepared for. The chapters that follow provide the framework for doing the work that comes next.

Chapter 1: The Evolution of the CISO Role

To understand where the CISO role is going, you must understand where it came from. The Chief Information Security Officer is a relatively recent invention in the annals of corporate leadership, and its evolution tells the story of how organizations have come to understand the value and vulnerability of digital assets.

From IT Police to Strategic Executive

The origins of the CISO trace back to the 1980s and 1990s, when information security was primarily the responsibility of IT operations teams. Early security professionals focused on keeping physical servers locked, managing passwords, and installing antivirus software. Security was viewed as a technical problem to be solved by technicians, not a business risk to be managed by executives.

The term "CISO" first appeared in the mid-1990s, but the role was rare and typically subordinate to the CIO. Security was a cost center, a compliance obligation, and a barrier to business velocity. Security professionals were often described as "the people who say no."

Several inflection points reshaped this perception. The Y2K crisis demonstrated the systemic risk that technology failures posed to global operations. The rise of organized cybercrime in the early 2000s, exemplified by worms like Code Red and Blaster, showed that attacks could spread at scale and cause real financial damage. Data breaches at companies like TJX, Heartland Payment Systems, and Target revealed that poor security practices could lead to multi-million dollar losses, regulatory penalties, and lasting reputational harm.

Sarbanes-Oxley, enacted in 2002, mandated internal controls over financial reporting that required security oversight of financial systems. GDPR, implemented in 2018, introduced heavy penalties for inadequate data protection and

elevated data security to a legal obligation. HIPAA, PCI-DSS, and other sector-specific regulations added further requirements. Each regulation pushed security from the IT department into the boardroom, because compliance failures now carried legal and financial consequences for the organization as a whole.

The 2010s marked a period of professionalization. CISOs began appearing on executive teams rather than nested under IT. Professional organizations like (ISC)2, ISACA, and SANS developed certifications and research that gave the discipline intellectual credibility. Analyst firms like Gartner, Forrester, and Ponemon began publishing annual benchmarks and studies that allowed CISOs to measure themselves against peers.

By the 2020s, the role had undergone a fundamental shift. The SolarWinds supply chain attack in 2020 demonstrated that even organizations with mature security programs could be compromised through trusted third parties. Ransomware attacks on Colonial Pipeline, MGM Resorts, and Change Healthcare showed that cyber incidents could disrupt critical infrastructure and entire industries. The pandemic accelerated cloud adoption and remote work, expanding the attack surface beyond any organization's perimeter. And now, artificial intelligence has introduced capabilities and threats that did not exist five years ago.

In this environment, the CISO is no longer just responsible for keeping systems secure. The CISO is responsible for ensuring that the organization can operate securely in a hostile digital environment. That is a fundamentally different mandate.

The Rise of Enterprise Risk and the Modern CISO

The modern CISO must operate at the intersection of technology and enterprise risk. Enterprise risk management, which encompasses financial, operational, strategic, compliance, and reputational risk, has increasingly recognized cyber risk as one of its most significant components. Yet cyber risk differs from traditional risk categories in important ways.

Financial risk can be quantified with historical data and statistical models. Operational risk can be managed through redundancy and process controls. Strategic risk can be mitigated through diversification and competitive analysis. Cyber risk, by contrast, involves intelligent adversaries who adapt,

learn, and innovate. The probability and impact of cyber incidents cannot be predicted with the same precision as market volatility or equipment failure. And the consequences of a major incident can span financial loss, regulatory penalties, customer attrition, operational disruption, and reputational damage, often simultaneously.

This complexity requires a different approach. The modern CISO must translate technical risk into business language that executives and boards can use to make decisions. That means moving beyond metrics like number of vulnerabilities or phishing click rates and instead answering questions like: What is the financial exposure if a specific risk materializes? Which risks matter most to our business objectives? How much should we invest in mitigation versus accepting residual risk?

Frameworks like FAIR, the Factor Analysis of Information Risk, provide quantitative methods for expressing cyber risk in financial terms. FAIR breaks risk down into two core components: frequency of loss events and magnitude of loss. By analyzing scenarios and aggregating outcomes, organizations can estimate potential loss exposure in dollars rather than abstract risk scores. This enables comparison between cyber risks and other business risks, supporting better investment decisions. [1]

The modern CISO also recognizes that prevention alone is insufficient. Defense in depth is valuable, but no organization can prevent all attacks. The focus has shifted toward resilience: the ability to detect incidents quickly, respond effectively, and recover rapidly. This requires capabilities across the incident lifecycle, not just at the perimeter.

A key indicator of maturity is how the CISO frames their role internally. Immature programs position security as compliance enforcement: checking boxes to satisfy regulations and audit requirements. Mature programs position security as business enablement: reducing risk to support strategic objectives, protecting customer trust, and ensuring operational continuity. The most mature programs integrate security into business decisions at the point where those decisions are made, rather than reviewing them after the fact.

Reporting Lines That Matter: Board, CEO, CIO, and CRO

The question of where the CISO should report has been debated for years, and it remains one of the most consequential structural decisions an organization makes about its security program.

The traditional model has the CISO reporting to the CIO. This makes some logical sense: security is closely related to IT operations, and the CIO has budget authority over the infrastructure that needs protecting. But this structure also creates inherent conflicts of interest. The CIO is responsible for delivering business value through technology, often measured by speed, innovation, and cost efficiency. Security requirements can slow delivery, increase costs, and complicate architecture. When the CISO reports to the CIO, security priorities may be deprioritized in favor of delivery timelines.

The alternative model has the CISO reporting to the CEO. This signals that security is a board-level concern, not an IT issue. It gives the CISO direct access to the executive who is ultimately responsible for organizational risk. It also provides independence when security and delivery priorities conflict. However, this model is not appropriate for every organization. In smaller companies, the CIO may be a more effective partner for building security capabilities. In organizations where the CIO is mature and security-forward, reporting to the CIO can work well if the CISO also has regular, direct access to the board.

Another emerging model has the CISO reporting to the Chief Risk Officer. This integrates cyber risk into enterprise risk management and provides a natural partner for risk quantification and regulatory reporting. But CROs often lack technical depth, which can make it harder for them to evaluate the CISO's recommendations or advocate for security investments with the CEO and board.

There is no universally correct answer. A CSO Online article notes that "ideally, the CISO would report to the CEO or the general counsel, high-level roles explicitly accountable for enterprise risk. Security is fundamentally a risk and governance function, not a cost-center function." [2] The National Association of Corporate Directors recommends that boards review whether the CISO's reporting line remains fit for purpose, particularly when regulatory requirements increase or when a material security incident has occurred. [3]

The most important factor is not the title or the box on the org chart but the actual access the CISO has to decision makers. A CISO reporting to the CIO who is invited to board meetings when cyber issues are discussed and has direct access to the CEO may be better positioned than a CISO reporting to the CEO who is isolated from technical and business stakeholders. Whatever the structure, the CISO needs three things: authority over security decisions, visibility to the board, and access to the budget required to execute the security

strategy.

What Changed: Data, Cloud, Ransomware, and Regulatory Pressure

Four forces have fundamentally reshaped the CISO mandate in the past decade.

First, data has become the primary asset. In industrial organizations, the most valuable assets were physical: factories, machinery, inventory. In digital organizations, the most valuable assets are data: customer information, intellectual property, proprietary algorithms, transaction records. Protecting data requires different skills and tools than protecting infrastructure. Data must be classified, discovered, encrypted, monitored, and governed across its lifecycle. And data protection is no longer solely a security issue. It intersects with privacy, legal, compliance, and business strategy.

Second, cloud computing has eliminated the traditional perimeter. Organizations that once controlled their own networks now operate across multiple cloud providers, with workloads distributed globally. The shared responsibility model means that cloud providers secure the infrastructure, but customers are responsible for securing their configurations, access controls, and applications. Yet many organizations struggle with this division of responsibility. A Check Point report found that 82 percent of enterprises experienced security incidents due to cloud misconfigurations. [4] The CISO must now govern security across environments they do not physically control.

Third, ransomware has evolved from an annoying nuisance to an existential threat. Early ransomware encrypted files and demanded payment to restore access. Modern ransomware groups exfiltrate data before encryption, threatening to publish it if ransoms are not paid. They target supply chains to amplify impact. They specialize in industries, with dedicated teams for healthcare, financial services, and critical infrastructure. The FBI's Internet Crime Complaint Center received 3,156 ransomware complaints in 2024 with adjusted losses of over 12.4 million dollars. [5] The average ransom payment in 2024 reached 2.73 million dollars, nearly double the previous year. [6] CISOs must now prepare for ransomware as a baseline threat, with tested backups, incident response playbooks, and business continuity plans that assume disruption will occur.

Fourth, regulatory pressure has intensified. The GDPR introduced fines of up to four percent of global annual revenue for serious violations. The CCPA and similar state laws in the United States gave consumers rights over their data and required organizations to respond to data requests. DORA, the Digital Operational Resilience Act, imposes stringent requirements on financial entities in the European Union, including mandatory resilience testing and third-party risk oversight. The SEC requires public companies to disclose material cybersecurity incidents within four business days and describe their cybersecurity risk management processes in annual reports. And the EU AI Act introduces new obligations for organizations deploying high-risk AI systems. [7]

Each regulation adds complexity, cost, and accountability. The CISO must coordinate with legal, privacy, compliance, and business teams to navigate this environment. And in some jurisdictions, regulatory action can target individual executives for failures in oversight, elevating cyber risk from an organizational concern to a personal liability.

The Current State: Where CISOs Succeed and Where They Fail

The current state of the CISO profession reflects both significant progress and persistent challenges.

On the positive side, security has gained strategic prominence. Deloitte's survey of 1,200 decision-makers across 43 countries found that 73 percent saw increased CISO involvement in strategic technology conversations over the prior twelve months. [8] CISOs are increasingly participating in product strategy discussions, M&A due diligence, and board-level risk governance. Organizations recognize that security is a competitive advantage: customers and partners demand strong security practices, and breaches can destroy trust built over decades.

Investment has grown. Security budgets as a percentage of IT spending rose from 8.6 percent in 2020 to 13.2 percent in 2024, according to IANS Research. [9] Organizations are investing in identity and access management, security operations centers, cloud security posture management, and AI-driven detection tools. Some are establishing dedicated product security, threat intelligence, and vulnerability management teams.

Frameworks and standards have matured. NIST CSF 2.0, ISO 27001, CIS Controls, MITRE ATT&CK, and other frameworks provide structured approaches to building and measuring security programs. These frameworks help CISOs articulate their posture to executives, regulators, and auditors.

But the challenges are real. First, CISOs face a talent crisis. ISC2's 2025 Workforce Study found that skills shortages persist even as staff shortages show slight improvement, with AI security (41 percent) and cloud security (36 percent) among the top deficits. [10] Organizations are competing for a limited talent pool while also facing budget constraints: 36 percent report budget cuts, 39 percent face hiring freezes. [11] CISOs must do more with less, while the threat environment grows more complex.

Second, CISOs face unrealistic expectations. Proofpoint's 2025 report found that 76 percent of CISOs expect a material cyberattack within a year, but 58 percent feel unprepared. [12] Executives and boards often expect security to prevent all incidents, yet they fund programs at a level that makes prevention impossible. This gap between expectation and reality contributes to stress and burnout. Proofpoint reports that 63 percent of CISOs experienced or witnessed burnout in the past year. [13]

Third, CISOs struggle with organizational alignment. The same Proofpoint survey found that board alignment dropped from 84 percent to 64 percent year-over-year. [14] When executives and boards do not fully appreciate the nature of cyber risk, they may underfund security programs, deprioritize security in business decisions, or blame the CISO for incidents that were inevitable given the resources provided.

Fourth, CISOs must navigate an expanding scope of responsibility. IANS Research found that up to a quarter of CISOs have assumed responsibility for areas beyond traditional cybersecurity, including IT operations, data governance, AI oversight, and digital transformation. [15] While this can elevate the role, it can also dilute focus and stretch resources thin.

The CISOs who succeed in this environment share common characteristics. They are strategic, not just technical. They build relationships with stakeholders across the organization. They communicate risk in business terms, not technical jargon. They prioritize ruthlessly, focusing resources on the risks that matter most to the business. They embrace resilience as much as prevention. They invest in their teams, knowing that human capability is the most important security asset. And they accept that the job is fundamentally

about managing uncertainty and making decisions with imperfect information.

The chapters that follow will examine each of these capabilities in depth, providing frameworks, practices, and examples that can help CISOs navigate the challenges of the AI era.

Chapter 2: The AI Transformation and Its Security Implications

Artificial intelligence is transforming cybersecurity more profoundly than any technology since the internet itself. But the transformation is not one-dimensional. AI is simultaneously a weapon for attackers, a tool for defenders, a source of new vulnerabilities, and a driver of regulatory change. The CISO who understands this transformation, and who can articulate its implications to executives and boards, will be better positioned to lead than the CISO who treats AI as a buzzword or an afterthought.

AI Capabilities That Reshape Security: Offensive and Defensive

To assess AI's impact on security, you must first understand what AI systems can do. At its core, modern AI, particularly large language models and generative AI, excels at pattern recognition, content generation, and complex decision-making based on vast datasets. These capabilities have direct implications for both offensive and defensive security operations.

On the offensive side, AI lowers the barrier to sophisticated attacks. A moderately skilled threat actor can now use generative AI to produce phishing emails that are grammatically correct, culturally appropriate, and highly personalized. AI can analyze public data about targets, including organizational charts, recent news, and social media profiles, to craft messages that reference specific events, projects, or relationships. This level of personalization, which previously required significant manual research, can now be automated at scale.

AI can also generate malicious code. Generative models trained on vast repositories of source code can produce malware, exploits, and obfuscation scripts that evade signature-based detection. Researchers have demonstrated AI systems that generate polymorphic malware, variants that change their appearance while preserving their functionality, making them harder to detect

with traditional tools. AI can also assist in vulnerability discovery by analyzing code repositories to find patterns that may indicate security flaws.

On the social engineering front, AI's impact is already profound. Deepfake technology enables voice and video impersonation of executives and trusted individuals. The Arup incident, in which a finance employee was tricked by AI-generated video calls of senior executives into authorizing a 25.5 million dollar transfer, demonstrated that deepfakes are no longer a futuristic concern. Voice cloning requires only 20 to 30 seconds of source audio. Video deepfakes can be produced in under an hour. The quality has reached a point where humans detect deepfakes only about 55 to 60 percent of the time, and automated detection tools fail in 45 to 50 percent of real-world cases. [1]

On the defensive side, AI offers significant potential benefits. IBM's 2024 Cost of a Data Breach Report found that organizations using AI and automation in their security operations saved an average of 2.2 million dollars per breach and reduced the total breach lifecycle by 98 days. [2] AI can analyze vast volumes of security telemetry, identify patterns indicative of malicious activity, and generate alerts that would be impossible for human analysts to detect at scale. Machine learning models can establish baselines of normal behavior and flag anomalies that may indicate compromise. AI can assist with vulnerability prioritization by correlating technical severity with business context, helping security teams focus on the issues that matter most.

AI is also being used for threat intelligence, automating the collection and analysis of indicators from disparate sources. Generative AI can draft incident response playbooks, summarize technical reports for executive audiences, and assist with forensic analysis by correlating logs and evidence. Some organizations are experimenting with AI-powered threat hunting, using models to proactively search for signs of compromise that evade traditional detection.

The generative AI in cybersecurity market is projected to grow nearly tenfold between 2024 and 2034, according to market research. [3] But growth in investment does not guarantee effectiveness. AI tools in security operations must be validated, monitored, and integrated into processes. They must reduce false positives, not add noise. They must augment human analysts, not replace judgment where it matters most. And they must themselves be secured against manipulation and abuse.

How AI Lowers the Barrier to Sophisticated Attacks

The most immediate implication of AI for security leaders is that it makes sophisticated attacks accessible to actors who previously lacked the skills or resources to execute them. This phenomenon, sometimes called threat democratization, has significant consequences for the CISO mandate.

Consider phishing. For years, phishing campaigns relied on obvious indicators: poor grammar, suspicious domains, generic greetings, and urgent language. Users were trained to recognize these cues. AI-generated phishing eliminates many of these giveaways. The messages are grammatically correct, contextually relevant, and convincingly written in the target's language and cultural context. This increases click-through rates and credential compromise.

AI also enables automation at scale. A single threat actor can use AI to generate thousands of personalized phishing emails, analyze which ones succeed, and refine their approach based on the results. This creates a feedback loop that improves attack effectiveness over time. Social media, data breaches, and public records provide the raw material for personalization. AI tools can scrape this information, correlate it with targets, and generate tailored content without manual effort.

Malware development is another area where AI lowers barriers. Traditional malware authors needed deep programming skills and knowledge of operating systems, memory management, and anti-analysis techniques. AI-assisted development allows less skilled actors to generate working malware from natural language prompts. An attacker can ask an AI system to write ransomware that encrypts files in a specific directory, excludes certain file types, and communicates with a command-and-control server at a specified address. The AI produces working code that can be deployed with minimal modification.

AI can also assist with evasion. Machine learning models can generate obfuscation techniques that make malware harder to detect by static analysis tools. AI can analyze detection rules and suggest modifications to bypass them. And AI can create variants of known malware families, altering their code enough to evade signature-based detection while preserving their core functionality.

Perhaps most concerning is the use of AI in reconnaissance and planning.

AI tools can analyze an organization's public footprint: websites, job postings, vendor relationships, technology stack, organizational structure. From this information, AI can suggest attack vectors, identify potential vulnerabilities, and generate initial access payloads tailored to the target's environment. This level of intelligence, which previously required a dedicated threat intelligence team, is now available to any actor with access to AI tools.

The implications for the CISO are clear. You cannot assume that sophisticated attacks are only launched by sophisticated adversaries. You must assume that AI tools will be used against your organization, whether by state-sponsored threat actors or opportunistic criminals. Defense strategies must adapt accordingly, focusing on detection and response capabilities that work against AI-enhanced attacks, not just the legacy techniques that defined the past decade.

AI in Defense: Augmentation, Automation, and New Failure Modes

AI offers powerful capabilities for security defense, but CISOs must be realistic about both the benefits and the limitations.

AI is already transforming security operations centers. Automated triage tools can analyze alerts, correlate them with contextual data, and prioritize them for analyst review. This reduces alert fatigue, a chronic problem in security operations where analysts are overwhelmed by thousands of alerts daily, many of which are false positives. AI-powered tools can draft initial investigation notes, generate containment recommendations, and even execute automated response actions for known threat patterns.

Machine learning models are being used for user and entity behavior analytics, establishing baselines of normal activity for users, devices, and network segments. When behavior deviates from the baseline, the system generates an alert. This approach is effective for detecting insider threats, compromised accounts, and lateral movement by attackers. Unlike rule-based detection, which requires pre-defined signatures, machine learning can identify novel attack patterns that have not been seen before.

AI is also being used for vulnerability management, a domain where organizations are typically overwhelmed. An organization may have tens of thousands of vulnerabilities to manage across its environment. AI can help prioritize them

by analyzing factors like exploit availability, asset criticality, business impact, and threat actor interest. This enables security teams to focus remediation efforts on the vulnerabilities most likely to be exploited.

But AI in defense introduces new failure modes that must be understood. First, AI models can be deceived. Adversarial machine learning techniques can craft inputs that cause models to misclassify malicious activity as benign. Attackers who understand the AI models being used for detection can develop evasion strategies that exploit model weaknesses. This creates an arms race: as defenders deploy AI for detection, attackers develop AI-powered techniques to bypass it.

Second, AI models can produce false confidence. A model that appears accurate in testing may perform poorly in production, especially when faced with data that differs from its training set. Security teams must validate AI tools rigorously, monitor their performance continuously, and maintain human oversight for critical decisions. AI should augment judgment, not replace it.

Third, AI introduces its own vulnerabilities. AI systems consume data, execute code, and make decisions. If the underlying models or training data are compromised, the AI system can produce harmful outputs. Prompt injection attacks can manipulate AI assistants into revealing sensitive information or performing unintended actions. Model poisoning attacks can corrupt the behavior of AI systems by manipulating training data. These vulnerabilities require new security controls that many organizations are only beginning to develop.

Finally, overreliance on AI can erode human capability. If analysts become accustomed to AI-generated recommendations and automated responses, they may lose the skills needed to investigate complex incidents independently. A balanced approach uses AI to handle routine tasks and provide decision support, while investing in human development for the situations that require experience, judgment, and creativity.

The CISO's responsibility is to evaluate AI tools critically, deploy them where they provide genuine value, and manage the risks they introduce. This requires understanding not just the marketing claims of vendors but the actual capabilities and limitations of the technology.

The New Attack Surface: Models, Data Pipelines, Prompts, and Agents

As organizations deploy AI systems, they create new attack surfaces that did not exist before. The CISO must understand these surfaces and govern them with appropriate controls.

At the foundation is the model itself. Large language models and other AI systems are trained on vast datasets. If that training data contains vulnerabilities, biases, or malicious content, the model will learn them. Model theft, in which an attacker extracts the weights or architecture of a proprietary model, can expose intellectual property and competitive advantage. Model inversion attacks can reconstruct training data from model outputs, potentially revealing sensitive information.

Then there are prompts. Prompt injection is now recognized as the number one security vulnerability for LLM applications by OWASP. [4] In a prompt injection attack, an attacker crafts inputs that manipulate the model into ignoring its original instructions and following attacker commands instead. Direct prompt injection involves crafting prompts directly to the AI. Indirect prompt injection involves embedding malicious instructions in data that the AI ingests, such as web pages, documents, or email. When the AI processes that data, it executes the embedded instructions, potentially leading to data exfiltration, unauthorized actions, or system compromise.

Microsoft's security research has documented cases where prompt injection led to remote code execution through AI agent frameworks that had access to system tools. [5] When an AI agent can read files, run commands, and access APIs, a successful prompt injection is no longer just a content issue. It is an execution risk.

Data pipelines are another critical surface. AI models require data for training, fine-tuning, and inference. That data flows through complex pipelines involving ingestion, transformation, storage, and access. Each step in the pipeline is a potential attack point. Data poisoning attacks can introduce malicious data into training sets, corrupting the model's behavior. Access controls on training data must be as rigorous as controls on production data.

AI agents, which are systems that can perceive their environment and take autonomous actions, introduce the most significant new risks. Agents are designed to make decisions and execute tasks without constant human oversight.

They may have access to databases, APIs, file systems, and communication platforms. If an agent is compromised, or if its decisions are manipulated through prompt injection or other means, the impact can be severe.

A Cloud Security Alliance report found that AI agent security incidents are now common in enterprises, with issues ranging from agents accessing data they should not have to agents taking unintended actions based on manipulated inputs. [6] Gartner predicts that by 2028, 25 percent of enterprise breaches will trace back to AI agent abuse. [7]

The CISO must govern these new surfaces with clarity. That means maintaining an inventory of AI systems in use, understanding what data they access and what actions they can perform, implementing appropriate access controls and monitoring, and testing them for vulnerabilities using techniques like red teaming and adversarial testing.

Why Traditional Security Perimeters Are Failing

The combination of AI, cloud computing, remote work, and complex supply chains has rendered traditional security perimeters obsolete. The concept of a defined boundary between trusted internal networks and untrusted external networks no longer reflects reality. Organizations now operate across cloud providers, with workloads distributed globally, users accessing resources from anywhere, and data flowing through complex ecosystems of third parties.

AI amplifies this reality. AI systems ingest data from the internet, process it through cloud-based models, and generate outputs that are distributed across the organization. The attack surface is not the network edge. It is the entire ecosystem of systems, data, and identities that AI depends on.

Traditional perimeter-based defenses, such as firewalls and network segmentation, remain valuable tools. But they are insufficient on their own. Defense must be layered, with security controls applied at every level: identity, data, application, infrastructure, and network. Zero trust architectures, which assume no implicit trust and verify every access request, are designed for this reality. But zero trust implementation is complex and expensive, requiring investment in identity management, network architecture, monitoring, and culture change. [8]

The CISO must also recognize that humans are part of the perimeter. Social engineering attacks that exploit human psychology, now enhanced by

AI, bypass technical controls entirely. Security awareness training remains important but is often insufficient. Organizations must design processes, such as dual approval for financial transactions and verification protocols for sensitive communications, that reduce reliance on individual judgment.

Finally, the supply chain is part of the perimeter. When organizations depend on third-party software, services, and infrastructure, they inherit the security posture of those dependencies. The SolarWinds, MOVEit, 3CX, and XZ Utils attacks all demonstrated how a single compromised component can cascade through thousands of organizations. [9] The CISO must manage third-party risk not as a compliance exercise but as a core element of defense strategy.

The perimeters have dissolved. The new paradigm is defense in depth across a distributed, dynamic environment where trust is earned, verified, and continuously validated. That is the security mandate of the AI era, and the CISO must be prepared to execute it.

Chapter 3: The New Risk Landscape

Nation-State Actors, Criminal Enterprises, and Hacktivists

AI-Enabled Threats: Deepfakes, Automated Social Engineering, and Malware

Adversarial AI: Prompt Injection, Data Poisoning, Model Theft, and Evasion

Supply Chain, Open Source, and Concentration Risk

Geopolitical Fragmentation, Digital Sovereignty, and Jurisdictional Risk

Chapter 4: Building an AI-Era Cybersecurity Strategy

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Starting with Business Outcomes, Not Threats

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Context Analysis: Industry, Regulation, Threat Intelligence, and Maturity

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Strategic Priorities: Prevention, Detection, Response, and Resilience

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

The Strategy Document: Structure, Content, and Executive Readiness

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

From Strategy to Roadmap: Sequencing, Dependencies, and Investment

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Chapter 5: Security Governance: Board, Accountability, and Decision Rights

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Board Expectations: What Directors Actually Need to Know

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Governance Structures: Committees, Charters, and Escalation Paths

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Decision Rights: What the CISO Owns, What Is Shared, What Is Delegated

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Regulatory Accountability and Personal Liability

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Making Governance Work: Meeting Cadence, Agendas, and Follow-Through

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Chapter 6: Enterprise Risk Management and Business Alignment

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

From Cyber Risk to Enterprise Risk: Integration Frameworks

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Risk Quantification: FAIR, CVSS, and Business-Impact Models

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Risk Appetite Statements and Board-Level Thresholds

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

The CISO vs. CRO: Overlap, Tension, and Collaboration

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Communicating Risk: Scenarios, Uncertainty, and Decision Quality

Chapter 7: Designing the Modern Security Organization

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Org Design Principles: Scale, Complexity, and Centralization vs. Embedding

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Common Security Organization Models and Trade-Offs

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Building the Security Operations Center: People, Process, Technology

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Outsourcing, Managed Services, and the MSSP Relationship

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Embedding Security: DevSecOps, Product Security, and Business Unit Liaisons

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Chapter 8: Talent, Culture, and Leadership

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

The Talent Crisis: Scarcity, Competition, and Retention Challenges

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Skills for the AI Era: What CISOs and Teams Must Know

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Leadership Models: From Heroics to System Thinking

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Building Security Culture Beyond Awareness Training

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Managing Burnout, Incident Stress, and Continuous Pressure

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Chapter 9: Budgeting and the Economics of Cybersecurity

Budgeting Models: Baseline, Risk-Driven, and Portfolio Approaches

Building the Business Case: From Technical Need to Financial Justification

ROI, ROSI, and Other Economic Metrics: What Works, What Does Not

Prioritization Frameworks for Scarce Resources

Cyber Insurance: Risk Transfer, Underwriting, and Strategic Implications

Chapter 10: Identity, Zero Trust, and Access Governance

Why Identity Is the New Perimeter

Zero Trust Architecture: Principles, Phases, and Implementation Reality

Privileged Access Management: The Critical Crown Jewels

Identity Governance: Lifecycle, Segregation of Duties, and Compliance

Balancing Security and Experience: Friction as a Strategic Choice

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Chapter 11: Cloud Security and Secure Architecture

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Shared Responsibility, Shared Confusion

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Cloud Security Posture Management and Configuration Baselines

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Platform Engineering and the CISO Partnership

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Container Security, Serverless, and Infrastructure as Code

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Secure Architecture Reviews: Process, Guardrails, and Velocity

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Chapter 12: Data Security, Privacy, and Regulatory Alignment

Data as the New Crown Jewels: Classification, Discovery, and Protection

Encryption, Tokenization, and Key Management at Scale

The CISO and the CPO: Collaboration, Tension, and Clarity

Regulatory Landscape: GDPR, CCPA, DORA, AI Act, and Sector Rules

Breach Notification and Cross-Border Data Flows

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Chapter 13: Application Security and Software Supply Chain

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

The Broken Window: Why Application Security Remains a Problem

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Shift Left Without Dogma: Practical Integration into Development

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Software Bill of Materials and Supply Chain Transparency

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Open Source Risk: Dependencies, Licenses, and Vulnerability Management

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Secure SDLC: Processes, Tooling, and Organizational Behavior

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Chapter 14: Security Operations and Incident Response in the AI Age

Detection in the Noise Age: Signals, Context, and Prioritization

AI-Assisted SOC Operations: Gains, Risks, and Validation

Incident Response Playbooks: Design, Rehearsal, and Real-World Execution

Threat Hunting and Proactive Defense

Post-Incident Analysis: Learning, Improvement, and Accountability

Chapter 15: Securing AI Systems: Models, Data, and Infrastructure

The AI Attack Surface: Models, Pipelines, APIs, and Data

Data Poisoning and Integrity: Prevention and Detection

Adversarial Attacks: Evasion, Membership Inference, and Model Extraction

Secure AI Development: MLOps Security and Model Governance

Monitoring AI Systems in Production: Drift, Abuse, and Anomaly

Chapter 16: Managing AI Adoption Risks Across the Enterprise

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Shadow AI: The Problem of Unmanaged Generative AI Use

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Data Leakage and Sensitive Information in AI Platforms

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

AI Usage Policies: Principles, Controls, and Enforcement

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Approved Tools, Sandboxing, and Enterprise AI Gateways

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Enabling Innovation While Managing Risk

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Chapter 17: Third-Party Risk, Ecosystem Security, and Concentration

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

The Third-Party Risk Explosion: Why Vendors Matter More Than Ever

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Vendor Risk Assessment: Processes, Questionnaires, and Reality Checks

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Concentration Risk and Single Points of Failure

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Contractual Protections, SLAs, and Breach Obligations

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Building Ecosystem Resilience: Shared Standards and Collaboration

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Chapter 18: Communication, Influence, and Executive Storytelling

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Board Communication: Cadence, Content, and Decision Support

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Executives Speak Business: Translating Risk into Dollars and Outcomes

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Influencing Without Authority: The Political CISO

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Crisis Communication: Breaches, Media, Regulators, and Customers

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Building Trust: Transparency, Honesty, and Consistency

Chapter 19: Resilience, Business Continuity, and the CISO Mandate

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Beyond Prevention: Resilience as the New Baseline

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Business Continuity and Disaster Recovery: Shared Ownership

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Ransomware Preparedness: Backups, Segmentation, and Playbooks

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Operating Under Attack: Graceful Degradation and Decision Making

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Resilience Testing: Exercises, Tabletops, and Realistic Scenarios

Chapter 20: The Future CISO: 2030 and Beyond

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Autonomous AI Agents: Security Implications for Organizations

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

AI-Driven Security Operations: What Changes for the CISO?

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Quantum Computing and Cryptographic Risk

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Regulatory Trajectories: Personal Liability, Mandatory Reporting, and Standards

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Geopolitical Conflict and National Cyber Posture

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

The Enduring CISO: What Will Not Change

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

Conclusion: The Mandate

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

The Non-Negotiables: What Every Modern CISO Must Own

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

The CISO as Chief Trust Officer

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

A Letter to the Next Generation of Security Leaders

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.

References

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/theai-eraciso.