THE CHIEF AI OFFICER

How to Build, Lead, and Scale AI Across the Enterprise

ERIC J. WHITMAN

The Chief AI Officer

How to Build, Lead, and Scale AI Across the Enterprise

Steve Publications

This book is available at https://leanpub.com/thechiefaiofficer

This version was published on 2026-09-08

This is a Leanpub book. Leanpub empowers authors and publishers with the Lean Publishing process. Lean Publishing is the act of publishing an in-progress ebook using lightweight tools and many iterations to get reader feedback, pivot until you have the right book and build traction once you do.

Contents

How to Build, Lead, and Scale AI Across the Enterprise

Artificial intelligence is no longer a technology experiment. It is a board-level strategic imperative. This book provides CEOs, board members, and senior executives with the authoritative guidance they need to establish, govern, and scale the Chief AI Officer role, the organizational structures that support it, and the AI capabilities that deliver real business value. You will learn when your organization truly needs a CAIO, how the role fits within the C-suite, what authority and accountability it requires, and how to build an AI operating model that turns strategy into measurable financial results while managing risk. The book covers everything from AI vision and maturity assessment to governance, security, vendor strategy, team building, financial management, and organizational transformation, grounded in research, real-world case studies, and practical frameworks you can apply immediately.

Introduction

The first time I saw a real boardroom AI failure, it was not because the technology did not work.

A mid-sized financial services company had spent three years and tens of millions of dollars building machine learning models for credit risk assessment. The models were accurate on test data. They were peer-reviewed. They passed internal validation. They sat in production, connected to the right systems, monitored by the right dashboards.

And nobody used them.

The credit officers still made decisions based on spreadsheets, experience, and email conversations with relationship managers. The models generated insights that arrived too late, or in formats nobody trusted, or that contradicted decisions already made by someone further up the organization chart. The company had solved the technical problem and completely missed the organizational one.

That scenario repeats across industries, now amplified by generative AI. Organizations rush to build pilots, hire data scientists, contract with vendors, and deploy models. Then the pilots expire. The vendors renew contracts. The data scientists move on. And the board asks the question it asks every year: so what did AI actually do for us.

This book exists to help executives answer that question differently.

Artificial intelligence has become a general-purpose technology, one that can reshape how an organization competes, operates, and creates value across every function. That level of impact requires a level of leadership that most organizations have not yet built. The Chief AI Officer emerged as the answer to that gap, not because the technology demanded it, but because strategy, risk, and transformation demanded it.

The CAIO is not simply a senior technology leader with a new title. The role bridges domains that traditionally live in silos: strategy and operations, data and analytics, product and technology, security and compliance, finance and human resources, ethics and innovation. A CAIO must speak the language of

the boardroom and the engineering floor, translate technical possibility into financial consequence, and govern risk without killing momentum. That is not an easy role to design or to fill. It is also not a role every organization needs.

In the pages that follow, you will find a complete framework for thinking about the CAIO role and the enterprise AI capability that surrounds it. The book is organized into twenty chapters, building from foundation to execution:

Chapters 1 through 4 define the CAIO role, trace the evolution of enterprise AI, and provide frameworks for building AI vision, assessing maturity, and designing roadmaps. If your question is should we hire a Chief AI Officer or where do we even start, these chapters provide the answer.

Chapters 5 through 9 cover the operating model, team building, technology foundations, vendor strategy, and data strategy. These chapters answer how do we organize, build, and fund this capability so it actually delivers.

Chapters 10 through 13 address governance, regulation, security, and lifecycle management. These are not peripheral topics. They are the guardrails that separate sustainable AI programs from programs that create legal, financial, or reputational risk faster than they create value.

Chapters 14 through 17 cover measurement, financial management, organizational transformation, and industry-specific impacts. If your challenge is proving AI ROI, managing budgets, driving adoption, or understanding how AI changes your specific function, these chapters are your reference.

Chapters 18 through 20 analyze real-world case studies, map strategic tradeoffs, and equip the CAIO to lead at the board level. These chapters show what exceptional AI leadership looks like in practice, including examples of failure that every organization should learn from.

A few guiding principles:

First, this book avoids hype. Artificial intelligence is not going to save your organization, and it is not going to destroy it. It is a set of capabilities that can produce extraordinary results when applied with clarity, discipline, and organizational alignment. It can also waste vast amounts of money and create serious risk when treated as magic.

Second, technology is not the core challenge. The hardest problems in enterprise AI are human and organizational: defining the right problems, aligning incentives, managing risk, building trust, driving adoption, and measuring

outcomes. Technical issues will appear, of course, but they are downstream of leadership decisions.

Third, there is no single correct answer. Different organizations need different structures, different strategies, and different degrees of centralization or decentralization. The goal of this book is to give you enough understanding to make the right decision for your organization, not to prescribe a template that fits everyone.

If you are a CEO or board member, this book will help you understand what to ask, what to expect, and how to hold leadership accountable for AI outcomes. If you are a CAIO, aspiring CAIO, or executive with AI responsibilities, it will provide frameworks, processes, and real-world evidence to guide your work. If you are a consultant or advisor, it will help you structure advice that is grounded in strategy and organizational reality rather than technology trends.

The AI advantage does not belong to the organization with the smartest models. It belongs to the organization that understands how to turn AI into a durable part of its strategy, operations, and culture. The Chief AI Officer is the role designed to make that happen. The question is whether your organization is ready to give it a real chance to succeed.

Chapter 1: The Emergence of the Chief AI Officer

Artificial intelligence did not become a board-level concern because the technology got smarter. It became a board-level concern because the stakes got higher.

When machine learning first entered the enterprise, it was housed inside IT, analytics, or research labs. Models predicted equipment failures, detected fraud, recommended products. The work was important but contained. The risks were managed. The impact was incremental.

Generative AI changed that equation. Large language models, diffusion models, and now autonomous agents can reshape how every knowledge worker in the organization does their job. They can improve customer experience or destroy it. They can generate revenue or hallucinate commitments that create legal liability. They can automate routine work or expose systemic bias in hiring, lending, and insurance decisions.

Impact at that scale cannot be managed by a manager inside IT. It requires an executive with a mandate that spans strategy, operations, technology, risk, and organizational change. That executive is the Chief AI Officer.

Why AI Requires a Distinct Executive Function

The argument for a dedicated AI executive is straightforward: artificial intelligence is too cross-cutting, too strategic, and too risky to be treated as a subfunction of IT or analytics.

Most enterprise AI programs span multiple business units, touch dozens of existing systems, create new regulatory exposure, and require changes in how employees work. That combination of breadth and complexity creates friction in any organization. When AI responsibility lives inside IT, the initiative tends to be viewed as a technology project rather than a strategic imperative. When it lives in analytics, it often lacks the product and operations mindset needed

to integrate AI into actual workflows. When it lives in research, it produces excellent papers and zero measurable business impact.

A Chief AI Officer exists to transcend those silos. The role is designed to connect AI capabilities directly to business strategy, ensure AI investments are funded and prioritized appropriately, govern risk across the portfolio, and drive organizational adoption. The CAIO is accountable for value creation, not just model performance.

The emergence of agentic AI systems has only strengthened the case for the role. Autonomous AI agents that execute tasks, interact with enterprise systems, and coordinate with each other create operational and security risks that require executive-level oversight. A report by Deloitte analysis of the MIT AI Risk Database identified over 350 risks from autonomous agent behavior, including runaway actions, misaligned goals, data exposure, and cascading errors in multi-agent systems[1]. Managing that level of risk requires someone whose job is to think about AI consequences at the enterprise level, not the project level.

From Data Science Manager to C-Suite Role

The Chief AI Officer did not appear overnight. The role evolved over more than a decade of enterprise AI experimentation, failure, and gradual professionalization.

The earliest references to a Chief AI Officer date to the early 2010s, when industry observers predicted that the CIO role would eventually evolve into an AI-focused leadership position. In 2016, Andrew Ng, one of the world's most prominent AI researchers, wrote in Harvard Business Review that organizations should consider hiring their first Chief AI Officer, arguing that companies treating AI as a sub-function of IT would fall behind competitors who elevated it to a strategic priority[2].

For a while, nothing much happened. Most organizations that used AI in production did so through internal data science teams or external vendors, with oversight distributed across IT, product, and risk management. A few large technology companies created dedicated AI organizations, but they were outliers.

The conversation accelerated during the 2020s as generative AI entered mainstream awareness. The rapid development of large language models,

the explosive growth of AI capabilities, and the visible productivity impact in knowledge work forced executives to rethink how AI was organized and governed. By 2024, a Gartner poll found that 54 percent of organizations had appointed an AI leader, though only 12 percent gave that leader the title of Chief AI Officer[3].

That gap between AI leadership and the CAIO title is important. Many organizations have created the function without the label: a senior vice president of AI and machine learning, a head of generative AI, a chief digital and AI officer, or a chief data and analytics officer who expanded their mandate to include AI strategy. The title matters less than the mandate. What matters is whether someone in the organization has clear responsibility for AI strategy, governance, and value creation at the executive level.

A Gartner survey of chief data and analytics officers found that 70 percent now have primary responsibility for building the AI strategy and operating model for their organizations[4]. That is consistent with an emerging trend toward combined roles: Chief Data, Analytics, and AI Officers who unify responsibility for the full data and intelligence stack. In a 2025 Harvard Business Review article, Vipin Gopal, Thomas H. Davenport, and Randy Bean argued that a combined CDAIO role would best prepare organizations for the AI era by unifying responsibility for data quality, analytical capability, and AI deployment under a single executive[5].

The trend toward combined roles makes sense in many contexts. Data, analytics, and AI are deeply interconnected. An organization cannot build effective AI capabilities without clean, accessible data. It cannot govern AI risk without understanding where data comes from and how it is used. Unifying these functions under a single leader can reduce organizational friction and improve coherence.

But there are limits. As AI capabilities grow more powerful and more visible to customers, boards, and regulators, treating AI as just another data capability can understate its strategic importance. Some organizations may eventually find that the CAIO role needs to stand on its own, reporting directly to the CEO or board, with authority that extends beyond the data and analytics function into product strategy, customer experience, operations, and risk management.

The CAIO Versus the CIO, CTO, CDO, CISO, and COO

Defining the Chief AI Officer role requires distinguishing it from other C-suite positions. In many organizations, AI responsibilities currently live with existing executives, and the question of whether to create a separate CAIO role depends on understanding where those existing responsibilities end and where the CAIO mandate begins.

The Chief Information Officer manages the internal IT estate: infrastructure, enterprise software, internal applications, end-user computing, and technology operations. The CIO ensures that the organization's technology systems are reliable, secure, and cost-effective. AI initiatives often depend on CIO-managed infrastructure and systems, but AI strategy is typically not the CIO's primary responsibility. The CIO's focus on stability and efficiency can sometimes work against the experimentation and speed that AI initiatives require.

The Chief Technology Officer is responsible for technology innovation, external technology partnerships, and the technical architecture of the company's products and services. The CTO is often the right home for AI when AI is primarily a product differentiator or when the organization is building AI-powered products for external customers. But when AI also transforms internal operations, customer service, risk management, and compliance, the CTO mandate may not extend far enough.

The Chief Data Officer owns data strategy, data quality, data governance, and data architecture. The CDO is a natural home for AI in organizations where AI is treated as an extension of analytics and business intelligence. Indeed, many organizations have expanded the CDO role into a CDAO or CDAIO. But as AI becomes more autonomous, more customer-facing, and more visible to regulators, treating AI as primarily a data function can understate its operational and strategic implications.

The Chief Information Security Officer is responsible for protecting the organization from cyber threats, data breaches, and security incidents. AI security is a subset of the CISO's broader mandate, but AI introduces unique security risks that require specialized expertise: prompt injection, model poisoning, data leakage through AI applications, and adversarial attacks on AI systems. The CAIO should work closely with the CISO to integrate AI security into the broader security program, but the CAIO also owns the business and

operational dimensions of AI risk that extend beyond traditional cybersecurity.

The Chief Operating Officer oversees business operations, process optimization, and execution of the company's operational strategy. AI is increasingly viewed as an operational capability, a tool for improving efficiency, reducing costs, and automating routine work. The COO may take ownership of AI initiatives that primarily affect operations. But when AI also affects product strategy, customer experience, innovation, and competitive positioning, the COO mandate may be too narrow.

The relationship between the CAIO and other executives is not zero-sum. A successful Chief AI Officer collaborates with every C-suite leader. The CAIO depends on the CIO for infrastructure, the CTO for product integration, the CDO for data quality, the CISO for security, and the COO for operational deployment. The CAIO role is designed to orchestrate those dependencies, not to duplicate or replace existing executive functions.

When AI responsibilities are distributed across multiple executives without a central owner, the result is typically fragmentation: duplicate investments, inconsistent governance, misaligned priorities, and unclear accountability. A CAIO provides a single point of accountability for the enterprise AI agenda, even while AI capability remains distributed across the organization.

When Your Organization Actually Needs a CAIO

Not every organization needs a Chief AI Officer. Many organizations can and should manage AI responsibilities through existing executives, augmented by new processes and structures. Creating a CAIO role when it is not needed can be counterproductive, adding a layer of management without adding clarity, creating turf wars with existing leaders, or setting expectations that the organization cannot meet.

Russell Reynolds Associates, a leading executive search firm, published guidance in 2024 on when organizations should and should not have an AI leader in the C-suite[6]. Their framework distinguishes between different levels of AI ambition:

For organizations where AI is primarily an enabling technology for improving productivity across departments, an AI leader embedded in technology or operations is often sufficient. For organizations where AI is critical to customer experience or product innovation, AI leadership should sit within the

commercial or product function. For organizations pursuing broad, holistic AI strategies that transform multiple business capabilities, a C-suite AI leader with enterprise-wide authority is appropriate.

A useful diagnostic is to ask five questions:

First, is AI central to the organization's competitive strategy? If AI is not a core driver of competitive advantage, a dedicated CAIO may be premature. If AI is fundamental to how the organization competes, a dedicated executive makes sense.

Second, is AI impact cross-functional? If AI will primarily affect one business unit or one function, that unit or function should lead. If AI will transform multiple functions, a centralized leader helps coordinate.

Third, is the risk profile significant? If AI decisions affect customer outcomes, regulatory compliance, safety, or financial exposure, a dedicated executive with governance authority is appropriate.

Fourth, is the investment substantial? If AI investments are material relative to the organization's budget, the level of investment warrants executive-level ownership and accountability.

Fifth, is the organization technically ready? If data infrastructure, data quality, and basic analytics capability are not in place, a CAIO cannot succeed. The organization may need a Chief Data Officer first, or a period of foundational investment, before adding a CAIO.

Helios HR, an executive advisory firm, found that 67 percent of mid-market organizations are in the foundational stage of AI adoption, exploring informally with no structured approach[7]. For organizations at that stage, a senior dedicated AI hire would likely outpace the work available, creating retention risk and unclear mandates. The right approach at that stage is to build foundational capability first: improve data quality, establish basic governance, run pilots that produce real value, and demonstrate the business case for dedicated leadership.

The wrong time to hire a CAIO is when the real problems are technical maturity, data infrastructure, or governance discipline. If the organization is still building foundational machine learning capabilities, fixing data quality, or establishing basic analytics practices, adding a C-suite AI leader will not fix those problems. It will only create a leader whose mandate cannot be fulfilled.

Reporting Lines and Where the CAIO Sits

When an organization does create a Chief AI Officer role, reporting structure matters. The CAIO's ability to influence strategy, secure funding, govern risk, and drive adoption depends heavily on where the role sits in the organization chart.

There are three common reporting models:

Reporting to the CEO. The CAIO sits at the same level as the CIO, CTO, CFO, and other C-suite executives. This structure provides maximum authority and visibility. The CAIO can directly influence enterprise strategy, participate in executive decision-making, and advocate for AI investments without filtering through another executive. This model is appropriate when AI is central to the organization's competitive strategy, when AI impact is cross-functional, and when the board expects direct visibility into AI governance.

Reporting to the CIO or CTO. The CAIO is housed inside the technology or IT organization, one level below the CEO. This structure makes sense when AI is primarily viewed as a technology initiative or when the organization is still building technical AI capability. It can reduce friction with existing technology leadership. But it can also limit the CAIO's influence on non-technology executives and make it harder to elevate AI to a true strategic priority. A CAIO reporting to the CIO may struggle to influence product strategy, customer experience, or organizational change when those functions report to other executives.

Reporting to the CDO or CDAO. The CAIO is combined with or reports to the Chief Data and Analytics Officer. This structure makes sense when AI is treated as an extension of data and analytics, and when unifying data quality, analytics, and AI under a single leader reduces organizational friction. The combined CDAIO role has gained traction, with many organizations consolidating these functions. But the risk is that AI remains viewed as an analytical capability rather than a strategic transformation.

The choice depends on organizational context. Gartner's polling data suggests that most AI leaders currently report to existing executives rather than sitting as independent C-suite roles[3]. That is reasonable for many organizations, especially early in their AI journeys. But organizations where AI is mission-critical should consider giving the CAIO direct access to the CEO and board.

Another consideration is whether the CAIO should sit on the executive committee. Sitting on the exco signals that AI is a core strategic priority and ensures the CAIO participates in strategic decision-making. For organizations pursuing serious AI transformation, CAIO membership on the exco should be the default.

Authority, Accountability, and Success Metrics

A Chief AI Officer without real authority is worse than no Chief AI Officer at all. The role must come with clear decision rights, accountability for outcomes, and success metrics that align with business objectives.

The CAIO's authority should extend across at least six domains:

AI strategy and roadmap. The CAIO owns the enterprise AI vision, sets strategic direction, and defines the roadmap for AI investment and deployment. The CAIO decides which use cases get funded, which capabilities to build internally, and which to acquire from vendors.

Governance and risk management. The CAIO establishes AI governance frameworks, defines risk classification, approves high-risk AI deployments, and ensures compliance with regulatory requirements. The CAIO has authority to stop or pause AI initiatives that create unacceptable risk.

Technology and platform. The CAIO defines the enterprise AI technology architecture, selects foundation models and AI platforms, and establishes standards for AI development and deployment. The CAIO coordinates with the CIO and CTO on infrastructure but has final authority on AI-specific technology decisions.

Talent and capability. The CAIO is responsible for hiring and developing AI talent, building internal capability, and managing external partnerships and vendors. The CAIO controls the AI budget for talent, infrastructure, and external services.

Measurement and reporting. The CAIO defines how AI performance is measured, establishes reporting processes for executives and the board, and holds AI initiative teams accountable for delivering business outcomes.

Organizational change. The CAIO leads organizational transformation related to AI adoption, including workflow redesign, training, communication,

and cultural change. The CAIO works with HR, operations, and business unit leaders to ensure AI capabilities are actually used.

Accountability must match authority. The CAIO should be held accountable for:

Value creation. AI initiatives must deliver measurable business outcomes: revenue growth, cost reduction, risk mitigation, or customer experience improvement. The CAIO reports on AI value using financial and operational metrics, not technical metrics.

Risk management. The CAIO is accountable for AI-related incidents, including data breaches, model failures, bias and discrimination issues, and regulatory violations. The CAIO must demonstrate that AI governance is working.

Capability building. The CAIO is accountable for developing organizational AI capability: hiring and retaining talent, building internal expertise, and creating reusable assets that enable future AI initiatives.

Board and executive alignment. The CAIO is accountable for keeping the board and executive team informed about AI progress, risk, and opportunity. The CAIO ensures that AI is treated as a strategic priority, not a technology project.

Success metrics for the CAIO role should include both financial and strategic measures. Gartner recommends five outcome-based metrics to prove AI value: sales conversion rate, average labor cost per worker, time to value, collection efficiency index, and employee net promoter score[8]. Other useful metrics include:

AI contribution to revenue and margin, measured through attribution analysis and controlled experiments

AI contribution to cost reduction, measured through efficiency gains and automation impact

AI risk incidents, measured through the number and severity of AI-related failures

AI adoption rate, measured through active usage and workflow integration

AI capability maturity, measured through standardized maturity assessments

The CAIO role is difficult. It requires technical depth, business acumen, leadership presence, and political skill. It requires the ability to balance speed and safety, innovation and governance, experimentation and execution. Organizations that create this role without giving it real authority, real resources, and real accountability should not be surprised when it fails. But organizations that invest in the role thoughtfully, that match the mandate to their actual AI ambition, and that support the CAIO with the structures and resources needed to succeed can gain a real competitive advantage.

The question is not whether AI matters. The question is whether your organization has the leadership to make AI work.

Chapter 2: The Enterprise AI Landscape

The role of the Chief AI Officer is defined by the technology landscape. Understanding how enterprise AI has evolved, where it is headed, and what it means for organizational design is essential for anyone responsible for AI leadership.

This chapter traces the technological and organizational evolution of enterprise AI from traditional machine learning to generative AI to autonomous agents, and explains how each phase changed executive responsibilities, risk profiles, and the organizational structures required to manage AI at scale.

The Three Waves of Enterprise AI

Enterprise AI has moved through three distinct waves, each with different technology capabilities, different use cases, and different organizational implications. Understanding these waves helps explain why the CAIO role has emerged now, and why treating AI as a single, monolithic technology is a mistake.

Wave 1: Predictive and Analytical AI (2010–2019)

The first wave of enterprise AI was dominated by predictive machine learning and advanced analytics. Organizations used machine learning models to forecast demand, detect fraud, predict customer churn, optimize pricing, recommend products, and improve risk assessment. The technology was powerful but specialized. Each model was trained for a specific task on a specific dataset. Models required data scientists to design features, train algorithms, evaluate performance, and deploy code.

The organizational footprint of Wave 1 AI was limited. AI teams were small, often housed inside IT, analytics, or research departments. The models they built improved operations incrementally but did not fundamentally change how most employees worked. AI was viewed as a technical capability, not a business transformation.

Wave 1 produced real value. Financial services used AI for fraud detection and credit scoring. Retail used AI for demand forecasting and recommendations. Manufacturing used AI for predictive maintenance. Healthcare used AI for clinical decision support and diagnostic imaging. But the value was contained. AI touched specific processes but did not reshape the enterprise.

Wave 2: Generative AI and Foundation Models (2020–2024)

The second wave began with the emergence of large language models and generative AI. Models like GPT-4, Claude, and Gemini could understand and generate human language at a level that was unprecedented. They could write code, summarize documents, answer questions, draft emails, translate text, and generate creative content. They were general-purpose tools that could be applied to dozens of tasks across every function in the organization.

Generative AI fundamentally changed the economics and organization of AI. For the first time, AI could be deployed at the individual employee level, not just at the enterprise system level. A single AI assistant could help a customer service agent write responses faster, a developer write code more efficiently, a marketer create campaign copy, a lawyer review contracts, and a researcher summarize literature.

The impact was immediate and visible. An NBER field study of 5,179 customer-support agents found that a generative AI assistant increased issues resolved per hour by 14 percent on average and by 34 percent for novice agents[9]. Other studies found similar productivity gains in knowledge work across writing, coding, analysis, and creative tasks.

But Wave 2 also created new challenges. Generative AI introduced risks that did not exist in Wave 1: hallucinations, where models generate confident but false information; data leakage, where models inadvertently expose sensitive information; bias and discrimination in model outputs; and copyright and intellectual property concerns. The technology was easy to use but hard to govern.

Wave 2 forced executives to take AI seriously. AI was no longer a back-office analytics capability. It was affecting customer experience, employee productivity, regulatory compliance, and competitive positioning. The scope and visibility of AI expanded so rapidly that organizations scrambled to figure out how to organize, govern, and scale the technology.

Wave 3: Agentic AI and Autonomous Systems (2025–present)

The third wave is just beginning. AI agents are systems that do not just respond to prompts but take action. They can execute multi-step tasks, interact with enterprise systems, coordinate with each other, and operate with varying levels of autonomy. An AI agent for customer support can not only draft a response but check inventory, process a refund, update the CRM, and send a confirmation email. An AI agent for software engineering can not only write code but run tests, fix bugs, and merge changes.

Gartner predicts that by the end of 2026, 40 percent of enterprise applications will feature task-specific AI agents, up from less than 5 percent in 2025[10]. Other surveys suggest even higher adoption rates, with one finding that 79 percent of companies are already running some form of AI agent.

Agentic AI changes the nature of AI risk and governance. When an AI system merely generates text or recommendations, mistakes are typically low cost. When an AI agent executes financial transactions, modifies critical systems, or interacts with external parties, mistakes can be expensive and damaging. The MIT AI Risk Database, analyzed by Deloitte, identified over 350 risks associated with autonomous agent behavior, including runaway actions, misaligned goals, data exposure, opaque decisions, and cascading errors in multi-agent systems[11].

Wave 3 AI also demands stronger organizational structures. Managing autonomous AI systems requires governance, monitoring, and control mechanisms that go beyond what was needed for Wave 1 or Wave 2. The CAIO role becomes more critical as AI systems become more powerful, more autonomous, and more deeply embedded in business operations.

Generative AI, Foundation Models, and the Productivity Inflection

Foundation models are the underlying technology that powered Wave 2 and continue to drive Wave 3. A foundation model is a very large, general-purpose model trained on vast amounts of data that can be adapted to many different tasks through prompts, fine-tuning, or integration with external tools. Foundation models can process text, images, audio, code, and other modalities.

The emergence of foundation models changed the economics of AI capability. Before foundation models, every organization had to build custom models

for every task. Foundation models provide pre-built capability that can be adapted quickly and cheaply. That reduction in cost and time to value made AI accessible to far more organizations and far more use cases.

The productivity impact has been substantial but uneven. A 2026 McKinsey report found that 80 percent of organizations now use generative AI and 62 percent have tested agentic AI, yet 60 percent report no enterprise-wide impact on EBIT[12]. The gap between adoption and impact is a central challenge for AI leaders.

Why does the gap exist? Generative AI produces the clearest value in specific contexts: tasks that are repetitive, knowledge-intensive, and rule-based, where AI can accelerate work without introducing unacceptable risk. Back-office operations, document processing, customer support, code generation, and content creation are examples where AI has demonstrated strong ROI. But applying AI across the enterprise requires more than technology deployment. It requires workflow redesign, training, change management, and governance, and most organizations have not invested enough in those non-technical requirements.

A 2025 MIT study reviewed over 300 generative AI implementations and found that only 5 percent generated millions of dollars in value[13]. The study identified organizational factors, not technical factors, as the primary differentiator. Organizations that aligned AI to clear business objectives, invested in change management, and measured outcomes rigorously outperformed organizations that focused primarily on technology.

The productivity inflection point for enterprise AI has not yet arrived at scale. But it is coming. As organizations gain experience, build better processes, and integrate AI more deeply into workflows, the measurable impact on productivity and financial performance will grow. The CAIO's job is to accelerate that inflection.

Agents, Tool Use, and Autonomous Workflows

AI agents represent the next frontier for enterprise AI, and they present both opportunity and complexity. Understanding agent architectures and their implications is essential for CAIOs and technology leaders.

An AI agent is a system that uses a foundation model to perceive, reason, and act. Agents can access external tools, call APIs, query databases, browse

the web, and execute multi-step workflows. They can maintain context across interactions, remember previous actions, and coordinate with other agents.

The architecture of enterprise AI agents typically includes several components:

The foundation model, which provides the reasoning and language capability.

Tools and APIs, which allow the agent to interact with external systems, retrieve data, and execute actions.

Memory systems, which enable the agent to maintain context, remember past interactions, and learn from experience.

Orchestration layers, which coordinate multi-step workflows, manage tool calls, and handle errors.

Guardrails, which constrain agent behavior, enforce policies, and prevent unsafe actions.

Agent architectures vary in complexity. Simple agents execute single tasks: answer a question, summarize a document, generate a report. Complex agents execute multi-step workflows: research a topic, draft content, review for accuracy, publish to a CMS, and notify stakeholders. Multi-agent systems include multiple agents that coordinate with each other, each specializing in a different capability.

The business implications of agentic AI are significant. Agents can automate not just individual tasks but entire workflows. A research agent can analyze market data, identify trends, and generate a report. A customer service agent can diagnose issues, process refunds, and update customer records. A software development agent can write code, run tests, debug issues, and merge changes.

But agent autonomy creates new governance challenges. When agents act independently, mistakes propagate faster. A single misaligned agent can generate spam emails, trigger incorrect financial transactions, or expose sensitive data before a human notices. Multi-agent systems introduce emergent behavior: agents interacting with each other can produce outcomes that were not anticipated by their individual designers.

Managing agentic AI requires governance that goes beyond traditional model risk management. The CAIO must ensure that agents are properly constrained, monitored, and logged. Agents must have clear boundaries on

what they can and cannot do. Agents must be designed with human oversight mechanisms that allow intervention when something goes wrong.

A Deloitte report on agentic AI risks in banking recommends six strategies for managing the risk: extending AI risk frameworks for agents, default agent disclosure, transparent agent identification and logging, agent-specific zero-trust cybersecurity, clear accountability with guardian agents, and specialized oversight talent[11]. These strategies are applicable across industries, not just financial services.

How Technology Evolution Changes the CAIO Mandate

Each wave of AI has changed what the CAIO role must cover. The role that made sense in Wave 1 is not the role that makes sense in Wave 3. Understanding this evolution helps organizations design the CAIO function appropriately.

In Wave 1, the CAIO's mandate was primarily technical and analytical: build machine learning models, ensure data quality, deploy models to production, measure model performance. The risks were limited, the regulatory exposure was low, and the impact was incremental.

In Wave 2, the mandate expanded to include organization-wide strategy, governance, and change management. Generative AI affected every function, required employee training, created new security and privacy risks, and attracted regulatory scrutiny. The CAIO had to coordinate across IT, legal, compliance, security, HR, and business units.

In Wave 3, the mandate must include autonomous system governance, agent architecture, operational risk management, and potentially new regulatory compliance. As AI agents execute real actions in production systems, the CAIO must ensure that agent behavior is safe, auditable, and aligned with organizational values.

The evolution of the CAIO role mirrors the evolution of AI from a technical capability to a strategic, operational, and regulatory imperative. Organizations that view the CAIO as a technology leader will underinvest in the governance, risk management, and organizational change that modern AI requires.

Implications for Risk, Investment, and Organizational Design

The three waves of enterprise AI have specific implications for how organizations think about risk, investment, and organizational structure.

Risk has grown. Wave 1 AI created primarily performance risk: models that were inaccurate or unreliable. Wave 2 AI added output risk: hallucinations, bias, inappropriate content, data leakage. Wave 3 AI adds execution risk: agents that take real actions with real consequences. Each wave requires stronger governance and risk management.

Investment has shifted. Wave 1 AI investment was primarily in data infrastructure, data science talent, and model development. Wave 2 AI investment expanded to include foundation model access, API costs, employee training, and governance frameworks. Wave 3 AI investment will include agent infrastructure, tool integrations, monitoring systems, and specialized risk management capability.

Organizational design has become more complex. Wave 1 AI could be housed inside a single function. Wave 2 AI requires cross-functional coordination. Wave 3 AI may require dedicated agent governance teams, specialized operational oversight, and integration with existing operational risk management.

The Chief AI Officer must evolve with each wave. The CAIO who succeeded in building predictive models five years ago may not have the skills or mandate to govern autonomous agents five years from now. Continuous learning, organizational adaptation, and strategic renewal are essential.

The enterprise AI landscape is changing faster than most organizations can adapt. The organizations that will succeed are the ones that treat AI not as a technology trend but as a fundamental shift in how work gets done, how decisions are made, and how value is created. The CAIO is the executive responsible for leading that adaptation.

Chapter 3: Defining the AI Vision and Corporate Strategy

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

From Hype to Strategic Intent

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

Mapping AI Capabilities to Business Priorities

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

Identifying and Prioritizing High-Value Use Cases

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

Portfolio Thinking: Balancing Exploration and Exploitation

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

Building the Business Case and Communicating to the Board

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

Common Strategic Failures and How to Avoid Them

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

Chapter 4: Assessing AI Maturity and Building the Roadmap

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

Dimensions of AI Maturity

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

The AI Maturity Assessment Framework

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

Diagnosing Gaps Between Ambition and Capability

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

Designing the Roadmap: From Pilots to Scale

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

Phasing Investments and Quick Wins

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

Measuring Progress Without Vanity Metrics

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

Chapter 5: Designing the AI Operating Model

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

The Centralized, Decentralized, and Federated Debate

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

AI Centers of Excellence: When They Work and When They Do Not

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

Platform Teams and Product Teams

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

Decision Rights, Funding, and Portfolio Management

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

Governance Structures That Enable Speed

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

Cross-Functional Collaboration and Friction Points

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

Chapter 6: Building and Managing AI Teams

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

Core Roles and Responsibilities in an AI Organization

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

The AI Talent Market and Hiring Realities

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

Retention, Career Paths, and Compensation

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

Upskilling and Reskilling the Existing Workforce

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

Working With External Specialists and Vendors

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

Creating Effective Collaboration Between Business, Data, Engineering, and Governance

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

Chapter 7: The AI Technology Foundation

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

Data Platforms and Infrastructure for AI

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

Cloud, Hybrid, and On-Premises Architectures

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

Model Platforms, APIs, and the Model Layer

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

Inference Infrastructure, Vector Databases, and RAG

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

Agent Architectures and Orchestration

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

Observability, Evaluation, and Security Controls

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

Chapter 8: Foundation Models and Vendor Strategy

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

The Foundation Model Ecosystem

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

Proprietary Versus Open-Source Models

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

Evaluating Models for Enterprise Use

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

Vendor Concentration, Lock-In, and Supply Risk

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

Negotiating Contracts and Commercial Terms

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

Build, Buy, and Partner Decisions

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

Chapter 9: Data Strategy for the AI Era

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

Data as the Fuel for AI

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

Data Quality, Ownership, and Governance

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

Unstructured Data, Knowledge Management, and Document AI

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

Data Lineage, Cataloging, and Accessibility

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

Privacy, Synthesis, and Synthetic Data

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

Aligning Enterprise Data Architecture With AI Demands

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

Chapter 10: AI Governance, Responsible AI, and Risk Management

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

Why AI Governance Cannot Be Optional

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

Model Risk, Hallucinations, and Reliability

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

Bias, Fairness, and Discrimination

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

Explainability, Transparency, and Human Oversight

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

IP, Copyright, and Content Provenance

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

Governance That Moves At the Speed of Business

Chapter 11: Regulatory and Legal Considerations

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

The Accelerating AI Regulatory Environment

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

The EU AI Act and Global Regulatory Fragmentation

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

Sector-Specific Regulation: Finance, Healthcare, Government

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

Contractual Risk, Indemnification, and Liability

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

Building an Adaptable Compliance Operating System

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

Managing Regulatory Uncertainty

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

Chapter 12: AI Security

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

The Expanded AI Attack Surface

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

Prompt Injection, Jailbreaking, and Adversarial Attacks

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

Data Leakage, Privacy, and Confidentiality

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

Secure Model Integration and API Security

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

Supply-Chain Risk and Third-Party Dependencies

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

Red Teaming, Monitoring, and Incident Response

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

Chapter 13: AI Lifecycle Management

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

Beyond Traditional Software Development

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

Experimentation, Development, and Evaluation

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

Deployment and Model Rollout Strategies

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

Monitoring, Drift Detection, and Retraining

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

Versioning, Registry, and Retirement

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

MLOps, LLMOps, and Platform Responsibility

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

Chapter 14: Measuring AI Performance and Business Value

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

Beyond Accuracy: A Multidimensional Evaluation Framework

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

Reliability, Latency, Cost, Safety, and Adoption

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

Linking AI Metrics to Business Outcomes

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

Building an AI Measurement System

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

Creating Executive Dashboards and Board Reports

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

Avoiding Vanity Metrics and Self-Delusion

Chapter 15: Financial Management and Investment Strategy

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

Understanding AI Cost Structures

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

Budgeting and Funding Models

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

Total Cost of Ownership and ROI Estimation

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

Infrastructure Economics and Inference Costs

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

Portfolio-Level Investment Decisions

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

Communicating Financials to CFOs and Boards

Chapter 16: Organizational Transformation and Change Management

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

Why Most AI Transformations Fail to Stick

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

Executive Sponsorship and Cultural Alignment

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

Workflow Redesign and Job Augmentation

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

Communication, Training, and Enablement

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

Managing Resistance and Job Anxiety

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

Building an AI-First Culture

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

Chapter 17: AI Across Functions and Industries

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

Impact on Knowledge Work and Professional Services

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

AI in Software Engineering and Product Development

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

Sales, Marketing, and Customer Experience

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

Operations, Supply Chain, and Manufacturing

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

Finance, HR, Legal, and Internal Functions

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

Industry-Specific Patterns: Healthcare, Financial Services, Government

Chapter 18: Case Studies

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

Enterprise Success Story: DBS Bank

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

Government and Public Sector Transformation

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

Financial Services: Risk and Innovation

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

Healthcare: Clinical AI and Regulatory Navigation

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

Manufacturing and Supply Chain Modernization

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

Failed Initiatives: What Went Wrong and Why

Chapter 19: Strategic Tradeoffs and Decision Frameworks

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

Centralized Versus Federated AI

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

Build Versus Buy Versus Partner

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

Proprietary Versus Open Models

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

Automation Versus Augmentation

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

General-Purpose Versus Specialized Models

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

Experimentation Versus Production Scaling

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

Chapter 20: Leading AI at the Board Level

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

What Boards Need to Know About AI

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

Presenting AI Strategy to the Board

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

Governing AI at the Board Level

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

Reporting on AI Progress and Risk

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

Building Board Competence and Accountability

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

The Future of the CAIO Role

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

Conclusion

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

References

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