

AGENTIC INTELLIGENCE

STRATEGY AT THE SPEED OF DATA



Manish Sood

Venkat Venkatraman

Forbes | Books

PRAISE FOR *AGENTIC INTELLIGENCE*

“This book delivers a pragmatic yet powerful roadmap for navigating the shift to an agentic intelligence era, where humans and machines work side by side. Through ten clear, actionable steps, it shows how unified, connected data becomes the foundation for everything that follows. Applicable to technical experts, business leaders, and board members alike, it is a practical guide to unlocking new value, driving smarter decisions, and staying competitive. Get the data right, and the possibilities are expansive.”

—**Jo-ann dePass Olsovsky**

Former CIO, Salesforce & BNSF Railway; Board Member, Reltio

“A must-read for every executive thinking about the AI readiness of their organization. Manish and Venkat lay out the business case for building a data foundation first. With real-life examples from the world of business, this book brings to life the transformational impact of actively managing and leveraging data.”

—**Mihir Shah**

Former CIO & Head of Data, Fidelity Investments

“As an executive or a practitioner, this book will become the book you go to for deep understanding of data in general and more importantly in a rapidly evolving AI world. Hard to think of strategy

or execution in an AI world without the concepts and examples in this book to guide you.”

—**Mike Gregoire**

Former Chairman & CEO, CA Technologies; Founding Partner, Brighton Park Capital

“A thought-provoking book using examples that inspire the reader to envision how they can create real, and real-time value from linked data. We are all navigating the opportunity of Artificial and Agentic Intelligence, but what’s blocking us and how can we go faster? Sood and Venkatraman bring answers to life in a realistic and simplified way using great case studies (‘Before Dawn’ as an example) followed by strategy-to-implementation Rules. You will relate to these stories so much that this book will become a trusted reference and roadmap that you’ll share with your colleagues.”

—**Carla Stratfold**

Former VP, Amazon Web Services, 2011-2021

“Having served on Manish’s board at Reltio, I have seen firsthand his rare combination of deep technical mastery and thoughtful, visionary leadership. He possesses a unique ability to cut through trends and focus on the fundamental architectural shifts required for true innovation. His insights are essential for any leader looking to build a resilient, future-ready organization.”

—**Chetan Puttagunta**

General Partner, Benchmark

“In *Agentic Intelligence*, Sood and Venkatraman take the concept of data-as-an-asset to the next level: data as a coordination fabric for

the organization. Monthly or quarterly strategy decisions will be for companies that find themselves in the dustbin of history. But so will those relying on AI as just a tool. This book makes the clear case for strategy as a system fueled by data in motion and actionable AI. Replete with visionary examples and a clear set of prescriptive advice, the book should be required reading for business, data, and AI executives today, even though its insights will certainly stand the test of time.”

—**Douglas Laney**

Data, Analytics, and AI Strategist; Author of Infonomics and Data Juice

“I found *Agentic Intelligence* to be one of the most grounded and actionable frameworks I’ve encountered. Sood and Venkatraman move beyond AI hype to address the foundational question of how organizations must redesign their data architecture and governance to enable humans and machines to truly work as coagents. Required reading for any executive serious about making AI operational at scale.”

—**Shamim Mohammad, CPA**

EVP, Chief Information & Technology Officer, CarMax

“Manish puts a new spin on where the future is headed with *Agentic Intelligence*. Where most books talk about the future, *Agentic Intelligence* showcases a path forward into where and how organizations will transform their businesses with this technology. Sood reminds of why having the right foundation matters in AI success. This book is an absolute must-read in the Age of AI.”

—**R. “Ray” Wang**

CEO, Constellation Research; 2x Best-Selling Author

OceanofPDF.com

AGENTIC INTELLIGENCE

STRATEGY AT THE SPEED OF DATA

Manish Sood

Venkat Venkatraman

Forbes | Books

[OceanofPDF.com](https://oceanofpdf.com)

Copyright © 2026 by Manish Sood and Venkat Venkatraman.

All rights reserved. No part of this book may be used or reproduced in any manner whatsoever without prior written consent of the author, except as provided by the United States of America copyright law.

Published by Forbes Books, Charleston, South Carolina.
An imprint of Advantage Media Group.

Forbes Books is a registered trademark, and the Forbes Books colophon is a trademark of Forbes Media, LLC.

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

ISBN: 979-8-88750-856-6 (Hardcover)

ISBN: 979-8-88750-857-3 (eBook)

ISBN: 979-8-88750-858-0 (Audiobook)

Library of Congress Control Number: 2026908656

No part of this publication may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the publisher.

The author reserves all rights to this work. This work, in whole or in part, may not be used for the purpose of training, developing, or improving any machine learning or artificial intelligence models, including large language models (LLMs) or generative AI systems, without express written permission from the author or rights holder.

This custom publication is intended to provide accurate information and the opinions of the author in regard to the subject matter covered. It is sold with the understanding that the publisher, Forbes Books, is not engaged in rendering legal, financial, or professional services of any kind. If legal advice or other expert assistance is required, the reader is advised to seek the services of a competent professional.

Since 1917, Forbes has remained steadfast in its mission to serve as the defining voice of entrepreneurial capitalism. Forbes Books, launched in 2016 through a partnership with Advantage Media, furthers that aim by helping business and thought leaders bring their stories, passion, and knowledge to the forefront in custom books. Opinions expressed by Forbes Books authors are their own. To be considered for publication, please visit books.Forbes.com.

To my wife, Victoria, for encouraging me to take the plunge into the world of entrepreneurship. And to all my Reltio team members, customers, and partners who joined us on the journey, for their unwavering support and belief.

—Manish

To my students—across universities, executive programs, and engagements—who have taught me as much as I have taught them, continually sharpening my thinking as strategy and digital technologies coevolve. And to my wife, Meera, and our two daughters—Tara and Uma—for the many conversations, debates, and provocations that keep the human and societal stakes of the digital world firmly in view.

—Venkat

OceanofPDF.com

CONTENTS

PREFACE

BY MANISH SOOD

PROLOGUE

A DAY IN 2030

SECTION 1

STRATEGIC INTENT

RULE 1

**AGENTIC INTELLIGENCE REQUIRES THAT DATA
BE LINKED, NOT MERELY COLLECTED**

RULE 2

**DATA IN MOTION—NOT DATA AT REST—CREATES
AGENTIC ADVANTAGE**

RULE 3

**AGENTIC INTELLIGENCE REQUIRES CROSS-
FUNCTIONAL, CROSS-PARTNER DATA**

RULE 4

**TRUSTED DATA IS THE NEW CURRENCY OF
COAGENCY**

SECTION 2

INVEST TO MONETIZE

RULE 5

INVEST ACROSS HORIZONS TO BUILD AGENTIC INTELLIGENCE

RULE 6

MONETIZE BY EMPOWERING AGENTS TO ACT

RULE 7

INVEST IN AGENTIC LITERACY TO BUILD SHARED CONTEXT

SECTION 3

IMPLEMENT TO WIN

RULE 8

LET AGENTS LEARN IN THE LOOP, NOT AFTER THE FACT

RULE 9

HUMANS AND MACHINES MUST ACT WITH THE SAME CONTEXT

RULE 10

GOVERN AGENCY IN MOTION, NOT DATA AT REST

EPILOGUE

AGENTIC TRANSFORMATION AS THE NEXT ARC OF
STRATEGY

BY VENKAT VENKATRAMAN

RELTIO AS CUSTOMER ZERO

BY MANISH SOOD

ABOUT THE AUTHORS

ENDNOTES

OceanofPDF.com

PREFACE

BY MANISH SOOD

It was late 2010. I wasn't sleeping well.

Now, the sorts of things that usually keep people up at night are career challenges, relationship problems, or money issues. None of that, however, was responsible for my sleep deprivation.

No, what was troubling me ... was data.

Since the mid-twentieth century, computers had given us the ability to comprehensively track data electronically. Over the decades since then, thousands, if not millions, of databases had been created by companies and individuals. That wasn't the problem. The problem was that most of those databases had completely different programming, interfaces, and applications, which made it impossible for them to communicate with each other and exchange information.

At first, it wasn't a big deal. Everyone was just excited to have the ability to track data. How the data was collected and what format it was in were secondary concerns. But then, as more data became available—especially with the rise of the internet—data became critically important for companies in many areas: marketing, operations, finance, and other industry-specific functions. However, most of the time, different databases served each department. Increasingly, there was a need for these systems to *work together*, but legacy systems often obstructed that. The longer these systems stayed isolated, the greater the risk that a company would become less efficient, less productive, and less competitive.

I knew it was an incredibly complex problem to unify all the data from those systems, but I also knew there had to be a solution. That was why insomnia was my closest companion at the time. I couldn't get it out of my head, and I saw the only way forward was to start a company that would tackle this looming crisis head-on.

That company was, and is, Reltio. That conviction has only grown stronger over time, and SAP's recent announcement that it intends to acquire Reltio is a powerful validation of the vision that first kept me awake at night in 2010.

Of course, you can't start a company like Reltio without cash. Because it was hard to explain our purpose to investors—people's eyes sometimes glaze over after just hearing the word *data*—it was difficult to raise capital at first. So, we had to change our approach. We initiated conversations with potential customers facing the problem we aimed to solve. They understood what we were trying to do and how it would help them—and they bought in. Investors noticed the positive response we were getting, which allowed us to finally secure sufficient financial resources to give it a real shot. I felt as if my vision for the company was validated.

Today, we serve many large global companies across more than 140 countries, in industries including financial services, healthcare, life sciences, insurance, retail, and hospitality. I'm sharing this not to boast but to highlight the breadth and depth of my experience in data. The truth is that my insomnia paid off; we are recognized leaders in helping clients build from the ground up for future success. We assist companies in leveraging their data resources to succeed today while laying a foundation for leadership in a world where artificial intelligence (AI) is becoming widespread in ways that would have been unimaginable at the start of this decade. When I first began this journey, I often found the data team tucked away in a company's basement. Now, data has rightfully earned its place in boardrooms, as leadership recognizes its vital role in driving business success.

This brings me to this book. In many ways, SAP's decision to acquire Reltio affirms the central argument Venkat and I make in these pages: that in the age of AI, competitive advantage depends on trusted, unified, real-time data and the context it creates. As data has become a crucial component of business operations, there is a need for a book that explains how a company should proactively and strategically manage data. If you're worried you'll be overwhelmed with complex technical jargon, don't be. This is not a book for engineers, coders, data scientists, or algorithm experts. It's intended for business leaders, future business professionals, and anyone whose role is heavily influenced by data and AI. Equally important, it's for any college student, regardless of discipline, who aims to work in companies of any size, anywhere in the world. Our goal is to show you how your professional skills can be enhanced and amplified by the power of data and AI.

You might be wondering about the name of this book: *Agentic Intelligence*. What does that mean? Simply put, since AI became mainstream in late 2022 (yes, thanks to OpenAI's ChatGPT), we have been given a whole new set of tools to streamline or reimagine business processes. More is possible and can be delivered faster thanks to AI agents, autonomous software systems that can understand goals, plan, and take actions using tools and data to achieve them on behalf of a user. Agentic intelligence will differentiate winners and losers. Our book will show you the way to victory and inspire you to lead.

We chose the subtitle carefully; *Strategy at the Speed of Data* captures a simple yet uncomfortable reality. AI has generated enormous excitement, yet its impact at scale remains limited. Most organizations can cite pilots and proofs of concept, but far fewer can demonstrate sustained results. This is not a failure of AI; it is a mismatch between how fast data now moves and how slowly organizations are designed to act. The recent SAP announcement is one more signal that this is no longer a niche data management problem but a board-level strategic imperative. Data today flows in real time—from customers, operations, and markets—while strategy, decision rights, and

governance still run on industrial-age clocks. Intelligence arrives faster than action, insight outpaces execution, and value is left unrealized. When we refer to agents, we mean both machines and humans. AI systems can reason and act at machine speed; humans provide judgment, experience, and accountability. Advantage comes not from replacing one with the other but from coordinating them so both operate with the same context and on the same clock. When machines move faster than humans can understand or trust, organizations slow them down. When humans must intervene everywhere, machines never scale. Agentic intelligence emerges only when humans and machines are designed to sense, decide, learn, and act together in real time. Therefore, the strategy itself must change. In the Intelligence Age, strategy cannot be a static plan or a periodic review; it must be a living system that connects real-time data to real-time decisions. What companies must focus on now is not more data or more pilots but learning how to compete using *strategy at the speed of data*.

Starting a book is intimidating, and I am fortunate to have a highly talented collaborator. Professor Venkat Venkatraman is Professor Emeritus at the Boston University Questrom School of Business, the author of *The Digital Matrix* (2017; 2023), and coauthor of *Fusion Strategy* (2024). He has been working at the intersection of business and information technology since he began his academic career at MIT Sloan School of Management in 1985. He has been thinking, teaching, researching, and advising companies longer than I have been practicing. His academic research and thought leadership have been essential in shaping the ideas in the pages ahead. Therefore, this book combines academic insight and practical experience—melding the best of both worlds.

I am sure you will agree that data has advanced, moving from the basement to the boardroom. This highlights the importance of real-time intelligence for a business to make well-informed decisions and operate efficiently. We have also begun delegating tasks and responsibilities to machines—whether they are digital agents on the cloud or physical robots among us. They can and will

work smoothly only with unified data that is contextualized and precisely tailored to their operating context. Such agents will not act independently of humans. Therefore, *agentic intelligence is the modern equivalent of organizational intelligence*. Whether you want to understand the agentic revolution we are currently experiencing or aim to lead your company's transformation, where humans and machines work together as coagents, I believe you will find this book enlightening and, hopefully, inspiring. It doesn't matter what your role, educational background, or professional aspirations are; you are already part of the world of agentic intelligence. Your role in shaping and being shaped by AI is just beginning.

Welcome to the Age of Agentic Intelligence.

[OceanofPDF.com](https://oceanofpdf.com)

PROLOGUE

A DAY IN 2030

The future depends on what we do in the present.

—MAHATMA GANDHI

Let's engage in a bit of time travel.

By 2030, technology has continued to transform nearly every aspect of life. Batteries charge within minutes and last much longer. Decentralized renewable energy grids have changed the way power is generated worldwide. Robots have shifted from novelty to routine; self-driving vehicles are common on the streets; and blockchain, along with other decentralized technologies, is redefining financial transactions. In healthcare, breakthroughs in gene therapy and longevity science are extending both lifespan and healthspan.

Artificial intelligence (AI) has also become more powerful and widespread. AI agents are now crucial coworkers (much more than tools) in most companies' technology systems. However, the promise of AI has not fully lived up to the hype as a cure-all. In fact, many organizations are struggling more than ever to manage information flow, still weighed down by tangled networks of outdated, incompatible legacy systems and processes that clash with newer technologies.

One corporation that isn't? AirMed Health.

Before Dawn

At 6:58 a.m. on an unseasonably warm October morning in 2030, a small amber dot blinked on AirMed's New York regional dashboard. Then a second. Then six more, almost in unison.

They did not belong to a hospital ward or a lab. They came from homes across South Brooklyn—detections of a quiet interplay of signals that no clinician had asked for and no patient had deliberately offered.

A teenager's smartwatch caught a pattern of interrupted breathing.

A middle-aged woman's smart speaker picked up a cough signature it didn't like.

A pharmacy website registered a cluster of searches for a specific antiviral.

Three families reported higher-than-normal overnight fevers from their sleep-tracking rings.

For years, such signals had been background noise—data without a destination. But now they flowed into a living, learning system.

Within eleven minutes, Sentinel, AirMed's regional health-state agent, triangulated them. "Possible respiratory cluster emerging, South Brooklyn," it alerted. It proceeded with the message: "Recommend early testing kits, targeted messaging, and pharmacy coordination."

At 7:14 a.m., AirMed's New York medical director swiped open the alert on her tablet. She watched as Sentinel mapped a faint but coherent pattern that had been invisible just thirty minutes earlier.

"For years," she murmured to her chief resident, "we waited for patients to get worse before we knew what was happening." She shook her head. "Now the system whispers it before anyone feels sick."

By 7:30 a.m., mobile testing vans were rerouting. Home test kits were being assembled for courier pickup. Pharmacies received early instructions to rebalance inventory. Patients were still brushing their teeth, unaware that a quiet piece of agentic intelligence was already protecting them.

At 8:02 a.m., CEO Elena Vargas stood backstage at the Midtown Hilton watching the cluster map update in real time. She smiled—softly, with the humility of someone who had seen the opposite of this.

It wasn't perfection she saw; it was *progress*. How much her organization had advanced in five years. And she knew exactly what it had taken to get there. A stage manager touched her arm and whispered, "They're ready for you."

She stepped into the light.

Investor Day, 2030

The ballroom doors closed softly behind her. Rows of investors, analysts, and industry observers, along with her management team, looked up with anticipation.

Investor Day presentations usually began with numbers—quarterly growth, year-over-year comparisons, adjusted earnings. But Elena, in her fifth year as CEO, didn't care to begin with numbers. Numbers came later. Today, she started with a story.

"Good morning," she said, gripping the lectern as if steadying the room. "Before we talk about performance, forecasts, and capital plans, I'd like to begin with a story."

A ripple of surprise. She tapped the clicker. A single image appeared on the forty-foot screen.

The shoebox.

The slide was titled:

"The Shoebox That Stopped the System"

Every time Elena saw the photo of the box—its worn cardboard edges, its lid buckling under a pile of receipts, scribbled doctor’s notes, CDs with unmarked scans, and printouts stained with coffee—she felt the same combination of disbelief and embarrassment. The audience laughed. Even now, a decade later, the absurdity still stung.

“This,” she said, “is why I am standing here.”

And she began the story.

Five years earlier, a seventy-five-year-old patient named Henry arrived at a crowded AirMed clinic with his entire medical history stored in a battered Florsheim shoebox. He had moved frequently, seen dozens of specialists, and had no better system than this cardboard container. It was his comfort zone—his safety blanket.

The receptionist froze. The nurses whispered. The clinic halted for almost an hour.

They could not take Henry in—not because his condition wasn’t urgent, but because *they had no way to understand his past*. They couldn’t decipher his state.

“We’ll need to reschedule,” the medical services administrator told him. “It may take a week to sort through this.”

Henry apologized for needing care. For being there.

The receptionist took a photo, muttering, “People won’t believe this.” She was right.

Elena clicked, and the shoebox exploded online. It became a symbol for everything broken about modern healthcare. Not just AirMed’s healthcare—everyone’s. Data locked in silos. Systems that didn’t speak. Patients forced to stitch their own histories together.

“That shoebox,” Elena said, “wasn’t our failure. It was our *mirror*.”

She let the silence settle again.

“And it was the beginning of our transformation.”

She clicked to advance to the next slide:

“When Patients Carried the Burden”

A new slide appeared: photos of plastic folders, Ziploc bags filled with pill bottles, Apple Watches showing unread metrics, a USB stick taped to an insurance card.

Elena walked away from the lectern and addressed the room directly.

“For decades, healthcare expected patients to be the integrators of their own care.”

A murmur passed through the audience—people shifting in their seats, some remembering their own files and folders.

She shared a quote from an AirMed patient survey in 2024:

“I feel like I’m the EMR (Electronic Medical Record).”

Another:

“I wear a device that knows my sleep cycle, my stress level, and my arrhythmia patterns—but my doctor sees none of it.”

And another:

“It’s frustrating that every doctor visit starts from scratch.”

She looked out at the investors. She continued.

“Our patients were not passive recipients. They were doing the heavy lifting—dragging their data behind them like anchors underwater.” She paused. “And five years ago, we decided to lift it.”

She clicked again. On the screen, the following words appeared:

“We Weren’t Just Siloed—We Were Blind”

A stark diagram replaced the photos—AirMed’s pre-transformation digital infrastructure. Six EMRs, three billing systems, four pharmacy platforms, two

scheduling systems, and dozens of add-ons.

It looked like a bowl of spaghetti built by committee.

“I became CEO two weeks after the shoebox went viral,” Elena continued. “My first act was to ask every department head to show us where data fragmentation was hurting patient care and business performance.”

She paused. “I asked for the good, the bad, and the ugly. I received mostly the ugly.”

She described the meeting vividly—fifteen leaders in a windowless conference room, each equipped with examples they had long been too exhausted or too afraid to share.

The pharmacy director spoke first.

He explained how a medication shortage in Chicago had led to untreated patients, only for AirMed to discover later that millions of pills sat unused in warehouses two states away.

“We don’t have a supply problem,” he had said. “We have a blindness problem.”

Then the chief nursing officer.

She told the room her nurses spent nearly half their shift reconciling mismatched orders, tracking missing lab results, and resubmitting paperwork that had vanished between systems.

“They are drowning,” she said, her voice cracking. “Not from work—but from the *wrong* work.”

The CIO went last.

He displayed a slide showing the maze of systems AirMed had accumulated through mergers and quick fixes.

“We are flying blind,” he concluded. “Our data exists. Our insights exist. But we cannot find them, cannot trust them, and cannot act on them.” He repeated: “We are just flying blind.”

Elena let that memory wash over the room.

“That was the day I realized something fundamental,” she said. “We didn’t need more data. We needed more *intelligence*.”

Then she clicked, and the next slide appeared with the title:

“Building the AirMed Neural Network”

“This slide,” she said, “represents the moment everything changed.”

On the screen was not a traditional organizational chart with siloed databases under each department. It looked more like a neural map—nodes connected by shimmering arcs.

“Our data scientists told us something uncomfortable,” she said. She mimicked—somewhat successfully—their tone. “Stop collecting data. Start connecting it.”

Laughter, but also recognition.

She described the early days of the “neural network” project—a small, overworked team led by a young data scientist named Riya.

“Riya told me in our first meeting, ‘Healthcare doesn’t need more and better databases. It needs a living map of understanding.’”

The team built

- real-time ingestion pipelines,
- context graphs to represent patients dynamically,
- machine agents that could reason across modalities,
- trust layers that recorded data lineage, and
- adaptive workflows that updated themselves.

But the work wasn’t glamorous. It was necessary. Foundational and fundamental to who we are now. She recounted an argument between Riya’s team and a legacy IT manager.

“Do you know what he said? ‘We’ve used this system for thirty years. It works.’”

She smiled. “So does a 1908 Model T. But I wouldn’t drive it into traffic.

“There were many failures. A pediatrics pilot that misclassified half of its early alerts; a lab-integration project that took eighteen weeks longer than planned; a supply chain model that recommended sending winter inhalers to Florida. And a few others. I was prepared for them as long as we, as an organization, learned from each one of them.

“What mattered,” Elena said, “was not perfection. It was *learning*.

“And slowly, the system learned. The organization learned. The neural network wasn’t just technical; it was organizational and foundational.”

She looked around the room, glancing at the corner where her team sat. She could see some heads nodding as she advanced to the next slide:

“The CFO’s Leap”

This slide had a different graphic; it showed a handwritten note—now famous inside AirMed:

“Elena—hospitals make money. Servers don’t. —Mark”

The room laughed. Elena shook her head with affection. “That was our CFO. And he was not wrong.”

She described the financial meeting where she first proposed allocating more than half of the new capital expenditures (capex) to data infrastructure, algorithms, and workforce retraining to leverage emerging AI agents—even if it meant delaying the opening of new clinics.

She remembered the stunned silence, the incredulous glances, and the board members shifting in their chairs. One investor had said, “You’re telling us to pour money into something we can’t see, can’t depreciate, and can’t tour or show off on our brochures?”

She had replied, “Yes. Because our future depends on what we *can’t* see yet.”

Mark—understandably—had resisted for months. He knew capex; he lived for real estate projects; he knew site expansions inside and out.

But machine intelligence? Real-time data flows? Continuous learning loops? Investing in experimenting with AI agents? He hadn’t been convinced.

Until one night in 2026.

Mark had found himself in the ER, not as a CFO, but as a father. It was at a different healthcare system. His daughter had an asthma flare that came out of nowhere. But something new happened. An agent flagged her biomarker patterns from her inhaler. Another cross-checked air quality data. Another pulled her prior responses to medication. The clinician walked in already knowing the likely cause, predicted severity, and optimal treatment.

Mark later wrote in his journal: “For the first time in my life, I trusted a digital system and understood its power to cure.”

He was convinced. He became a believer. He signed off on the new capital plan the next day.

Elena remarked, “Changing our mindset was not easy; it’s still not easy.” She clicked again. The following slide appeared with three—but essential—words.

“Governance Across Boundaries”

“This,” Elena said, “was the hardest part.” The slide showed not technology, but people—board members, regulators, pharmacists, insurers—all interconnected.

“Healthcare is an ecosystem,” she said. “But our governance cultures were not.”

For decades, every entity—labs, pharmacies, insurers, regulators, hospitals—had acted as if it were an island. “We needed a continental mindset,” she

said. She recounted board meetings full of tension:

Board Member Chen: “Why should we open our system? We’ll lose competitive leverage.”

Board Member Alvarez: “No. We’ll gain it. Transparency builds trust.”

Board Member Jones: “Governance is not a constraint—it’s our advantage.”

Elena described negotiations with a national pharmacy chain. Their CEO at first refused to share data with AirMed.

Until AirMed shared everything first—lineage, uncertainty, error rates, the reasoning paths of its agents.

“We exposed our intelligence,” Elena said. “Not because we were naïve, but because we were confident.”

The pharmacy CEO changed his mind. Then the labs. Then the insurers.

“You didn’t just share data,” the pharmacy CEO told her. “You shared your logic. You shared your vision for the future.” That changed everything. Elena paused and clicked.

“Ecosystem or Nothing”

The slide showed a living map of partners flowing into a unified stream of understanding.

“Once trust took root,” Elena said, “the ecosystem started to breathe.”

She told the stories of an insurer who complained that AirMed’s new real-time eligibility checks were “too fast”—their systems couldn’t keep up. A lab director who called her personally, saying, “You’re showing me test anomalies I didn’t know existed. How are you doing this?” A regulator who thanked AirMed for detecting a rare adverse interaction weeks before national surveillance systems caught it.

“These were not technology wins,” she said. “They were *coordination wins*. *That’s how business wins today*.”

She shared statistics—lightly, conversationally—“Delays fell by 78 percent. Prior authorizations dropped from fourteen days to six hours. Drug diversion incidents dropped to near zero. Seasonal flu surges were predicted one to two weeks early.

“But the most important moment,” she said, “came from a junior resident.”

He had scribbled on a whiteboard: “Feels like we’re hearing the heartbeat.”

She walked across the podium to take a sip of water and clicked to advance to the next slide.

“Agents in Human Hands”

The slide showed a nurse leaning over a patient—not staring at a screen.

“When we talk about agentic intelligence,” Elena said, “people imagine machines replacing humans.”

She shook her head. “That is not our philosophy. Agents don’t replace people. They give people *back* their time, their dignity, their purpose.”

She told the story of Nurse Helena, who once spent two hours a day writing notes and tracking tasks. Now, Helena has agents who draft notes, track vitals, predict risks, and coordinate follow-up tasks before Helena even asks.

One day, recently, Helena told Elena, “My agents don’t make me less human. They give me space to be human.”

A pulmonologist said, “I have twelve hundred patients on my roster. My agent watches them all. I focus on the twelve who need me today.”

A social worker said, “For the first time, the system remembers what I remember.”

Agents became the connective tissue—each one learning from human feedback, each one improving AirMed’s collective intelligence.

“Agents didn’t make our workforce smaller,” Elena said. “They made our workforce *smarter*.”

She tapped the clicker again.

“The Culture of Truth”

A live dashboard appeared—including bed occupancy, drug inventories, predicted bottlenecks, and patient outcomes. Not static charts but a dashboard with live updates.

“This slide represents the deepest cultural transformation we achieved,” she said.

Before, every meeting began with a debate about whose numbers were correct. Now, everyone saw the same thing. Real-time data. Not static data, but data in motion.

She described the first meeting held after the new “truth room” went live. Two directors began to argue about ICU utilization trends—then realized the dashboard in front of them had already resolved the discrepancy. Agents were working in the background to resolve the issue.

Suddenly, there was nothing left to argue about. Only decisions to make. Only patients to serve. And serve efficiently and humanely.

A veteran physician told her, “I didn’t know how blind I was until I could see again.”

A hospital administrator said, “Honesty became easier when the system was honest first.”

And a younger manager said, “It feels like working in daylight.”

Elena looked at the screen.

“A culture of truth is not about data,” she said. “It’s about trust. That’s what I’ve worked so hard to instill in AirMed over the past five years.”

As she clicked, the next slide appeared:

“Patients as Coagents”

It showed a composite illustration combining multiple indicators, including sleep cycles from wearables, blood oxygen saturation patterns, glucose trends, smart nutrition logging, and more.

“Some people fear that AI will make healthcare less human,” she said. “In fact, it made our patients *partners*.” She told the story of Maria, a cancer survivor whose recurrence was detected earlier because her wearable flagged a subtle fatigue pattern.

Maria had written Elena a letter:

“The first time I had cancer, the treatment felt like it was designed for someone like me. The second time, it felt *meant for me*.”

For the first time, AirMed hadn’t treated patients as categories, clusters, or profiles. It had recognized them as individuals with unique trajectories—dynamic, evolving, and deeply personal.

“Eliminating the shoebox wasn’t enough,” Elena said. “We had to eliminate the burden patients carried—literally and emotionally.

“Patients now have their own health agents: agents who interpret data, provide insights, and work alongside clinicians. These agents continuously know every individual; more importantly, they are trusted by patients to protect their privacy and identity.

“Coagency,” she said, “is the future of healthcare.”

Then, she clicked one last time. The slide featured the AirMed logo and seven game-changing words:

“We Learned to Act in Real Time.”

The final slide returned to the morning’s vignette—the cluster in South Brooklyn Elena had eyed before taking the stage.

The audience understood now. They had lived the journey through Elena’s eyes.

“What you saw this morning,” she said, “was not the triumph of machines. It was the triumph of a system that finally learned to see, understand, and act.”

She stepped away from the lectern. And continued.

“In five years, we rebuilt AirMed’s architecture, retrained our workforce, rewrote our governance, reshaped our ecosystem, and redefined our purpose.”

She paused—deliberately.

“And we discovered that intelligence is not a model or an algorithm. Intelligence is the *coordination of insight and action*, across humans and machines, across partners and patients.”

She looked around the ballroom, taking in the faces of skeptics now leaning forward in their seats.

“Our transformation is not perfect,” she said. “But it is real. And it is human.”

She concluded with the words that would later become AirMed’s unofficial credo:

“We became intelligent not to replace people, but to empower them. To give every nurse time to care, every doctor clarity and responsibility to cure, every partner our trust, and every patient their dignity.”

She closed her clicker. “This is AirMed. This is who we’ve become. And we’re by no means finished—because this is the beginning of the Intelligence Age.”

The room rose in applause.

* * * *

You might not work in healthcare or carry a shoebox of medical records into a clinic, but the forces that transformed AirMed are already affecting your world. If you're a manager in retail, logistics, finance, manufacturing, or higher education, you've experienced the same fragmentation that once challenged AirMed—disconnected systems, blind spots between teams, and the feeling that your customers understand themselves better than your organization does. What AirMed achieved isn't a healthcare miracle but a glimpse of what happens when humans and machines learn to work together with real-time information and shared context. Every company has its own version of the shoebox; the question is whether leaders will face it and have the imagination and courage to build something better—something suited for the rapidly approaching future coming faster than we could have imagined just a few years ago.

And suppose you are a student wondering about your future, an intern uncertain about your prospects, a frontline worker questioning where you fit as AI accelerates, or a patient navigating your own care. In that case, you are already part of this story. Agentic intelligence is not about replacing you; it is about empowering you. It makes your judgment more valuable, not less. It expands your opportunities, not your anxieties. It provides you with tools that think alongside you, not for you. In a world where data flows freely and agents amplify human strengths, you are not a spectator. You are a coagent—a participant in shaping how intelligence is used, shared, and governed. AirMed's story is simply one way to see what that future looks like when we choose to build systems that make people more capable, more trusted, and more seen.

You might still see this story as just a business dream, but we want to offer some perspective: AirMed's story isn't purely fictional; it offers a glimpse into the future. As William Gibson said, "The future is already here, just unevenly distributed."

Every company and industry now faces the same choice that our fictional CEO, Elena Vargas, made. Do we continue relying on a patchwork network of data systems, many of which are as outdated as Pets.com and Ask Jeeves are today? Do we choose to evolve, or do we let ourselves be held back by cumbersome, incompatible data systems? Do we embrace agentic intelligence—systems that don't just respond to inputs but act autonomously with initiative and agency—or continue to struggle with friction that hampers productivity and progress?

This moment calls for leaders to answer this question:

Do we invest to become the pioneers of our future—or do we let ourselves be prisoners of the past?

Change is hard. It's often costly and faces opposition from those who are too committed to “how things have always been done.” However, businesses are recognizing that this change is inevitable and that the opportunity to adapt is quickly disappearing, as those who resist will fall behind. Their operations will slow, become more expensive, and carry greater risk; their customer relationships will weaken; and any competitive advantage they had will erode.

So, the question is not whether the future will arrive. It already has. The only question is: *How will you reinvent yourself to win in the future?*

* * * *

We wrote this book with one purpose: to help you lead your organization's transformation and succeed in a business environment where agentic intelligence is central. Those who haven't taken steps to manage their data effectively will risk significant financial losses. Poor data quality costs the US economy billions of dollars annually.¹ For individual companies, the impact can be more directly devastating—affecting their profitability, competitiveness, and reputation.

So how do you begin a journey like AirMed's—moving from fragmented data and frustrated teams to a world where humans and machines work together with clarity, confidence, and purpose? The truth is simple: Every organization, no matter its size or sector, can make this leap. But it requires a shift in how leaders think about data—not as a byproduct of business, but as the strategic logic on which agentic intelligence is built.

This book offers ten rules to guide that transition. They're not just abstract theories or Silicon Valley catchphrases. They're based on years of academic research, practical experience in some of the world's most complex organizations, and direct observation of what happens when companies see data as a strategic asset rather than a hidden cost center.

To help you navigate them, we've organized the rules into three tightly interwoven strands—three forces that, together, shape whether a company can thrive in the Intelligence Age.

One: Strategic Intent. This is where transformation begins. Strategic intent defines how your organization sees data—whether as scattered records to be stored, or as living connections that power human-machine collaboration. It is the lens through which you envision the role of agents, the value of interoperability, and the kind of organization you aspire to build. Without a clear intent, investments scatter, systems drift, and AI initiatives stall. With intent, everything aligns.

Two: Invest to Monetize. Intent is only meaningful if you fund it. This second strand focuses on how leaders allocate capital—toward architectures, platforms, and talent capable of generating real economic value. This is where the business case for data becomes clear: reducing costs, accelerating decisions, unlocking new revenue models, and enabling scalable, autonomous workflows. Many organizations experiment with pilot projects; few invest with conviction. This section shows you how—and why—those that do achieve disproportionate returns and emerge as winners.

Three: Implement to Win. Even the best strategy and the healthiest budget are ineffective unless they change how work is truly done. Execution in the Intelligence Age emphasizes coordinating workflows where agents—human and digital—operate with shared understanding, context, and trust. This final element demonstrates how to redesign processes, governance, and roles so agents boost human capacity rather than compete with it. Execution is where vision becomes lasting through profits and growth.

These three strands—Strategic Intent, Invest to Monetize, and Implement to Win—are not sequential steps but a braided system that connects strategy to success with speed. Each reinforces the others. Intent guides investment. Investment enables implementation. Effective execution sharpens and reframes intent. This is how organizations shift from a legacy logic of data silos and process workarounds to a modern logic of interconnected intelligence.



Strategic Intent

Invest to Monetize

Implement to Win

In the chapters ahead, we invite you to explore these ten rules with curiosity and ambition. Our aim is not to overwhelm you with technical jargon or promise magic through algorithms. Instead, we've written this book for the thoughtful business leader—someone who may not write code but understands that the future of competition will belong to those who see data not as waste but as fuel.

Francis Bacon famously said, “Knowledge is power.” In the Intelligence Age, knowledge is no longer just power—it is partnership. It is coagency. It is the foundation for a future in which humans and machines work together to generate more value than either could alone.

This book will help you unlock that power, develop that partnership, and guide your organization into a future where human and machine intelligence—working together—become your greatest strategic advantage.

Most people are very familiar with artificial intelligence—machines that perform tasks or imitate human thinking. In this book, however, we talk about an intelligence that’s more impactful: *agentic intelligence*. It’s not about replacing humans with smarter tools, but about combining human judgment with machine agency to create a new kind of advantage. The next era will not be won by algorithms alone or human intuition alone; it will be defined by how well organizations design this partnership—where humans and machines work together, learn together, and redefine what advantage really means.

Welcome—once again—to the Age of Agentic Intelligence.

OceanofPDF.com

SECTION 1

STRATEGIC INTENT

We shape our buildings; thereafter they shape us.

—SIR WINSTON CHURCHILL

There's a well-known fable about a group of blind men who encounter an elephant for the first time. Each man is asked to describe what he thinks the elephant looks like. Each touches a different part of the animal. The result?

- One grabs the trunk and insists the elephant is like a snake.
- Another person feels the leg and says it's like a tree.
- A third person touches the ear and says it's like a fan.
- Another touches the tusk and says it's like a spear.
- Still another feels the tail and claims it's like a rope.

Now, each of these men is partly right—but also completely wrong. None of them can accurately describe an elephant. They can only accurately describe *parts* of the elephant.

What's the moral of this story? That you can't really understand something if you can't see the full picture.

This elephant metaphor is often used to illustrate the risk of siloed data in organizations: Each team holds a piece, but without integration, no one can see the full picture of the “elephant.” This leads us to the first section of the book—strategy and its crucial role in business data.

If our five men had come up with a clear strategy—something as simple as sharing their individual experiences—they might have been able to develop a reasonable idea of what an elephant looks like. Because they didn't? Chaos. Everyone experienced their own reality, but it didn't reflect the full picture of the pachyderm.

That's what happens to a company without a unified view of strategy. Without integration, collaboration, or a broader view, everyone ends up with a distorted, incomplete understanding—and may even argue over who's right. For example, let's say several functions in a company only have access to certain aspects of customer data. The result? Different numbers, leading to different conclusions:

- Marketing studies the click-through rates and declares, "The customer is all about engagement!"
- Sales looks at purchase history and argues, "No, the customer is all about transactions."
- Finance focuses on revenue per account and insists, "The customer is all about profitability."
- Customer service examines complaint logs and counters, "The customer is all about problems."

Like the proverbial elephant, each department is right in its own way but also wrong in its own way. Because the data is siloed, no one can see the whole customer. Or the entire company for that matter. The result is confusion, conflicting strategies, and wasted effort. This is disastrous when humans are the decision-makers and actors. What happens when AI agents—machines with decision rights to act to serve customers—operate under such conditions?

Now, the question is: How did we arrive at a place where this could happen?

Over the years, organizations have established one system after another to collect specific data in particular areas, rarely considering whether these systems might need to communicate in the future to exchange information efficiently.

To some extent, this was inevitable. The speed of technological progress over the past few decades has been rapid. It was logical for businesses to adopt these new systems to build faster, more powerful data systems to gain great benefits. However, few considered the major drawback—the newer software generations couldn't easily communicate with older ones. It's like putting a baby boomer in a room with a millennial and expecting them to see eye to eye—they don't speak the same language, even if they both speak English, Spanish, or Hindi.

Over time, many companies have faced increasing friction because of incompatible systems and data across functions. Initially, this friction was likely dismissed as minor—just a small distraction—but as unified data became increasingly crucial for businesses, more obstacles emerged. Still, most companies hesitated to overhaul their entire data systems to ensure seamless integration. It seemed too overwhelming—and too expensive—a daunting challenge. Most leaders decided it was easier to create work-arounds, even if they were awkward and unwieldy.

Until those work-arounds were unworkable.

That point came with the COVID-19 pandemic, which sparked a data panic. It revealed and worsened the challenges of siloed data in ways companies could no longer ignore. The pandemic required real-time insights into supply chains, customer needs, workforce availability, and cash flow in a rapidly changing marketplace, as customer behavior evolved over time. When those insights weren't available, many dominoes fell:

- Organizations with fragmented customer data across marketing, sales, and service departments *struggled to adjust messaging or personalize outreach.*

- Retailers, manufacturers, and distributors discovered their supply chain data was scattered across regions and systems. Without a unified view, they *couldn't anticipate shortages, reroute inventory, or quickly identify alternative suppliers*.
- Remote work made collaboration more difficult. When data sat in departmental silos, employees couldn't easily share or access what they needed. This led to *duplication of effort, confusion, and reduced productivity* at a time when agility was most crucial.

On the other hand, organizations that had invested in data unification before 2020 outpaced siloed organizations. It's common to find studies reporting that companies with strong data practices are more likely to acquire new customers and be profitable.

We've seen it firsthand with our clients. One such company is Mercury Insurance.

Mercury Insurance faced a critical challenge because of its legacy systems. When the pandemic struck, fewer people drove, leading California to require insurance companies to offer customers partial premium refunds. As Mercury Insurance was the fourth-largest carrier in the state, it needed to inform its customers about these refunds through their physical addresses, phone numbers, or email addresses. However, the data team told the executives that they couldn't be sure the data was reliable because it was too scattered. They used multiple applications for different types of insurance policies, and some data was stored in a single system, while other data was kept in separate systems.

That highlighted other data and workflow issues. For example, if an agent was speaking with a car insurance customer on the phone and tried to sell them a home policy, the agent had to manually upload their auto policy to verify coverage—again, because they needed to access different databases. This took time and created friction. A knowledgeable person with access to data

across multiple systems could manage this. But how would AI agents handle this task? Not very effectively, without frustrating customers who'd naturally scream to speak to a human!

Not only that, but when a Mercury customer with multiple policies logged on to the company's customer portal, they could only view siloed information about their vehicle and home coverage. They might have two or three other policies with Mercury, such as personal umbrella or mechanical protection coverage, but still be unable to see them. This led to many unnecessary calls to Mercury agents from frustrated customers, who wondered what had happened to their other policies.² Mercury realized that, to move forward, they needed a solid strategic plan to unify their data and keep it unified. We—at Reltio—helped them do just that.

You've likely experienced frustrations with companies that still lack a unified data system. They make you scroll through generic websites, endure frustrating chatbot interactions, and sit through endless phone calls with agents who put you on hold for what feels like forever to look up your records—just to fix a small issue that should take seconds. One major cell phone provider we encountered could not find a way to change someone's customer ID, even though their last name was—and still is—spelled incorrectly.

Siloed data is to blame for most of this. As the COVID-19 pandemic showed, companies can adapt quickly when necessary. Yet many still operate as Mercury did back in 2020—stuck in the Industrial Age, even though that era has long passed. We're now in the Intelligence Age, where data drives everything. Making that transition can feel daunting, as if it requires a Stanford or MIT engineering degree. It doesn't. The first step is simply to understand the key ideas behind data unification, then to find the right people to help you implement them. If you're a business leader, that process should feel familiar—it's how transformation always begins.

Every meaningful transformation begins long before the first system is replaced, the first dataset is cleaned, or the first AI model is deployed. It begins

with intent—with the clarity of purpose that determines not only what a company builds but also *why* it builds it. If leadership gets the strategic intent wrong at the outset, everything that follows will feel misaligned: Investments scatter, pilots stall, teams become skeptical, and promising technologies turn into expensive experiments.

That is why the first section of this book focuses on four rules that form the blueprint for agentic intelligence. They are not technical steps; they are design principles for how an organization chooses to think, learn, and act in the Intelligence Age. They shape the company's "nervous system"—the way information flows, decisions are made, and humans and machines coordinate in real time. Without these principles, the later rules—on execution, investment, and monetization—never reach their potential. With them, everything else becomes possible.

Taken together, these principles form the strategic logic of agentic intelligence. They are not about data or technology per se; rather, they concern how an organization defines itself. Architecture gives the structure for awareness. Velocity infuses responsiveness. Interconnection broadens perspective. Trust anchors legitimacy. When these four elements align, a company begins to behave less like a machine and more like an organism—sensing, reasoning, adapting—through a network of human and digital agents working together.

This is the essence of strategic intent: a deliberate choice to design your organization for intelligence rather than inertia. And just as no company would dream of operating without clear strategies for manufacturing, marketing, or customer acquisition, no company in the Intelligence Age can succeed without a coherent strategy for data. The threads of modern business—sales, supply chains, operations, finance, service—are all woven from data. Strategic intent determines whether that fabric is strong or frayed.

As you read the first four rules, consider them not as abstract advice but as design decisions about the kind of enterprise you want to lead. Everything that

follows—every investment, workflow, human capability, and profit—will rest on this foundation.



Strategic Intent

Invest to Monetize

Implement to Win



RULE 1

AGENTIC INTELLIGENCE REQUIRES THAT DATA BE LINKED, NOT MERELY COLLECTED

AI agents will become the primary way we interact with computers in the future. They will be able to understand our needs and preferences and proactively help us with tasks and decision making.

—SATYA NADELLA, CEO OF MICROSOFT

The shoebox was not an outlier.

Let's explain. In the prologue, we created a fictional story about a seventy-five-year-old man who kept all his medical records in an old, battered shoebox. For him, it was the best way to keep all his information in one physical place. It might have seemed as silly to you as it did to the people working at our fictional healthcare provider, but honestly, it's not much better than the system we've relied on for over a hundred years—and that some still use today.

We're talking about filing cabinets—both physical ones and metaphorical ones.

Let's start with the physical ones. Back in the day, filing cabinets were just as revolutionary as today's advanced data storage. In the nineteenth century, before they were widely used, records were stored randomly in envelopes, desk drawers, or flat boxes. If you needed to find specific information, you might end up on an endless scavenger hunt unless you knew exactly where to find the document.

But then a man named Edwin Seibels changed everything.

In 1898, Seibels had the clever idea to create a vertical filing system—to organize that jumble of papers and store them upright in labeled folders within metal cabinets. If you're thinking Seibels became incredibly wealthy from his invention, think again—interestingly, his patent was denied because it was considered an idea, not an actual device.³

However, filing cabinets became popular because suddenly, documents could be easily indexed, filed, and retrieved. By the mid-twentieth century, they had become essential for managing business information. The only drawback? In large organizations, each department had its own collection of cabinets, managed independently. This created a tradition of siloed data, where marketing's records weren't easily accessible to finance, HR's data was invisible to factory operations, and so on.

In the second half of the twentieth century, with the rise of mainframes and personal computers, organizations began digitizing records. Databases and enterprise systems eventually replaced physical filing cabinets. What was harder to replace was the mindset those cabinets represented. Instead of using the new technology to unify all data across the entire business, departments again created their own separate digital “cabinets,” keeping their databases isolated from one another. Databases, warehouses, and enterprise resource planning (ERP) modules were designed to store data, not to facilitate flow. Initially, this

wasn't a major concern because we were still in an industrial era of predictable cycles and quarterly and yearly reviews.

Then came the data explosion of the early twenty-first century and the advent of cloud-based data storage. Email, cloud applications, and digital transactions exponentially increased data volume. But the established data silos persisted. Yes, companies now have marketing automation platforms, customer relationship management (CRM) systems, HR systems, and finance tools, but they still store their own data in their own systems and tools. And just like paper silos once slowed collaboration, digital silos now prevent leaders from seeing the whole customer, the whole business, the whole picture.

In this new era, complete, unified access to all relevant data is essential, not optional. Fortunately, in recent years, modern platforms such as data lakes, master data management (MDM),⁴ and customer data platforms have helped companies move beyond the filing-cabinet mentality. We at Reltio have played our part in supporting companies in this domain. We've seen firsthand that those who invested wisely have broken down data barriers, integrated fragmented data, and delivered a single, real-time view of information across functions. Those who didn't or couldn't are now facing mounting challenges to remain competitive as the world turns toward the Intelligence Age.

How Datagraphs Drive Neural Networks

Older companies were built around the filing-cabinet method because there was no alternative at the time. To catch up, they must not only rethink their entire approach to data but also migrate information from outdated legacy systems, which can be a struggle. For newer organizations, there's no need for such a complex transition *if* the goal of creating neural connectivity is pursued from the outset. From the start, they will establish optimal conditions to use all their data for multiple purposes.

One of the clearest real-world examples of this principle is Aadhaar,⁵ India's national digital identity system. Aadhaar provides residents with a unique twelve-digit ID linked to basic biometric and demographic data. Since its launch in 2010, it has been adopted at an extraordinary scale—covering virtually the entire adult population—and is now used to open bank accounts, receive government benefits, file taxes, enroll in school, and access a wide range of public and private services. The World Bank has described it as one of the most sophisticated identity programs globally.

For citizens, the difference was immediate. Before Aadhaar, proving identity meant carrying stacks of documents and navigating inconsistent bureaucratic requirements. After Aadhaar, identity became portable, verifiable, and instantly usable: Showing a single ID replaced years of paperwork and uncertainty. Life became simpler not because there was more data, but because identity worked consistently across systems.

Aadhaar's success rests on a critical design choice: Interoperability was built in from the start. The ID was designed to function as a shared reference point across services, agencies, and platforms. By contrast, identifiers such as the US Social Security number were never designed as digital credentials and remain difficult to integrate securely across systems. Retrofitting them is as much a political challenge as a technical one.

The lesson for organizations is straightforward. What India achieved at a national scale, companies must achieve internally: connecting fragmented systems through shared identifiers and common context. This is how trapped value is unlocked. But connectivity alone is not enough. Signals without structure create noise, not insight. What makes a system intelligent—biological or digital—is not just speed but the ability to recognize relationships and meaning across connections.

At its core, Aadhaar is a foundational civic achievement: a universal, portable, and verified identity layer spanning more than a billion people. It solves a deceptively hard problem at a national scale—establishing who

someone is, reliably and consistently, across institutions and over time. That clarity has enabled practical outcomes, such as direct benefit transfers, faster financial onboarding through the electronic Know Your Customer feature, and integration across India's public digital platforms.

What Aadhaar does *not* do is equally important. It does not reason, predict, personalize, or act. It is not an intelligent system, and it does not operate as an agent on behalf of individuals. Its role is simpler—and more powerful: It anchors every person to a clean, trusted identity within a broader digital ecosystem. It creates a stable reference point on which other systems can rely.

That anchoring function is precisely what makes Aadhaar so consequential for the next phase: agentic services. Agentic systems depend on knowing who they are acting for, what permissions apply, and how information connects across domains. Without a trusted identity backbone, coordination quickly becomes fragile—permissions blur, attribution breaks down, and trust erodes. Aadhaar provides that backbone. It enables consent-driven connections between identity and services—financial, educational, health, and civic—not by creating intelligence itself but by ensuring every connection is tied to a real, verified individual.

For businesses, start-ups, and public institutions, this distinction matters. Aadhaar is not the agentic layer—but it is the precondition for one. It allows agentic services to be built on a stable, governed foundation rather than on fragmented identifiers and improvised trust. Its achievement is not that it delivers agentic intelligence. It is that it makes agentic intelligence possible, as well as coherent, authorized, and scalable.

That is the foundation of agentic intelligence: not more data but shared context that allows systems—and the people who work with them—to act coherently.

This is accomplished through *datagraphs*.

Most of you probably intuitively understand what datagraphs can do, even if you're unfamiliar with the term or haven't seen the graphs. For example, when you log on to Netflix, you're flooded with personal recommendations that focus on movies, shows, actors, and genres you've previously shown interest in, based on your viewing history. These recommendations appear surprisingly accurate, not because Netflix has the largest film and show database, but because it understands how all its offerings are connected—and which ones match your taste. As a result, 80 percent of what we watch on the streaming service is driven by personalized recommendations on the home screen.⁶

The datagraph principle isn't just for online shopping, streaming, and other consumer platforms—it's also transforming science and industry. AstraZeneca constructs and deploys biomedical graphs that link genes, diseases, compounds, and trial outcomes, allowing AI systems to suggest drug targets or repurpose therapies more quickly than manual literature searches. Siemens uses a digital twin that connects the physical machine to its virtual model by linking computer-aided designs, Internet of Things (IoT) sensor data, location information, and maintenance logs, enabling engineers and AI agents to predict failures across turbines and factories. JPMorgan treats datasets as reusable products within a data mesh, enabling AI to navigate clients, transactions, and risks with ease.

So ... what is a datagraph? Simply put, it's a corporate mind map. Rather than treating data as isolated tables, a datagraph links entities—such as customers, products, contracts, and events—into a dynamic network. Let's step back and picture a room full of Edwin Seibels's pride and joy: filing cabinets filled with important information. Well, the only way to connect that information is with a stapler. In contrast, imagine a wall-sized mind map instead of those filing cabinets—a wall map that connects data across silos and makes it accessible to AI agents and human teams. Those agents and teams can then create datagraph models that are extremely valuable to a company.

Datagraphs are more than just abstract technical ideas; they already deliver business value today. Take Ant Group (formerly Ant Financial), the fintech giant that spun off from Alibaba. Its datagraph acts as a high-precision “digital brain” that maps trillions of real-time financial relationships to address the industry’s toughest data challenges. By feeding verified facts into its AI systems, Ant Group reduces “hallucinations,” ensuring every credit decision and fraud check is grounded in reality. This enables the company to instantly assess the creditworthiness of small businesses that traditional banks usually overlook. By organizing its data as a neural network, it can process these complex connections at a speed and scale unmatched by competitors.

Or look at Estonia, the small Baltic nation that showed how even governments can wire data like a brain. Like India, Estonia built a digital identity system based on the once-only principle: Citizens provide data once, and it automatically flows across agencies. When they apply for a license, file taxes, or visit a physician, the systems connect with AI agents that are already beginning to streamline interactions. Compare that to what Mercury Insurance faced during the pandemic, when they couldn’t get a full view of all of a customer’s policies. For citizens, Estonia’s system feels seamless. For businesses, it’s a reminder that even highly bureaucratic institutions can operate in real time if the architecture is appropriate.

AI struggles in environments built for filing cabinets. Most organizations still operate in this world—the familiar maze of “systems of record” and “systems of engagement” that store data in isolated pockets. Humans can sometimes navigate these silos through improvisation or institutional memory. AI agents cannot. When data is locked in separate systems, an AI agent has no way to assemble the full story behind a query, no way to locate all the relevant signals, and no way to understand the relationships that give information its meaning. It’s the digital equivalent of asking a pilot to land an airplane without an air traffic control system: The plane may be powerful, the pilot skilled, but without a unified, real-time picture of the sky, the risk is overwhelming. This is

why chatbots hallucinate. This is why autonomous agents crash into invisible walls inside enterprises.⁷ And it's why so many promising AI pilots never scale: *Without context, intelligence collapses.*

These failures aren't about models—they're about the architecture. When data doesn't flow in real time, AI agents behave like air traffic controllers watching yesterday's radar. They miss new signals, depend on outdated information, and make decisions that no human would trust. When information is fragmented, even the most advanced models remain narrow, fragile, and easily confused. Traditional systems are designed for specific tasks, not for shared situational awareness. But agentic intelligence requires the opposite: a neural-network-like architecture that connects data across people, processes, and platforms. In such a system, humans and machines share a common view of reality, allowing them to coordinate and decide together.

With this architecture established, everything begins to click into place. A customer service agent—whether human or digital—can access not only the last complaint but the entire relationship history. A logistics agent can identify disruptions early and collaborate with planners to reroute shipments before delays occur. A clinical assistant can integrate patient histories, medical literature, and population health signals to produce evidence-based recommendations. These are not just futuristic promises; they are the results when organizations shift from isolated filing cabinets to building a system similar to air traffic control for their data—one that offers every agent, human or machine, real-time and connected context to act confidently. That's exactly what AirMed achieved in our stylized fictional example and why its AI agents—and people—can finally learn, act, and decide with clarity. And that's also how we helped AbbVie.

How We Helped AbbVie

Let's examine the case study of one of our customers, AbbVie—a global biopharma company that spun off from Abbott Laboratories in 2013. It currently employs over fifty thousand people worldwide and ranks as the second-largest biopharma company after Pfizer.

You can appreciate how vital data is to a company like this. Not surprisingly, the company had been dependent on legacy systems for its data. These systems were fragmented and slow, with limited governance and low user confidence. Because there were multiple platforms, AbbVie's MDM system performed poorly and was difficult to use.

How severe was the data situation? In the words of Vivian Wu (associate director at AbbVie): "It took an hour to do just one record."⁸ That hour was spent solely on merging records to ensure consistent data across platforms. These legacy systems also caused cumbersome, inflexible search options that made it hard for users to find what they needed. Data was isolated across multiple domains with no centralized governance or management system.

Finally, the company leadership decided that AbbVie needed to seamlessly integrate data from healthcare providers and other organizations. Their brief to us at Reltio was as follows:

- Replace the legacy systems with modern data management.
- Create new consumption and integration hubs, along with a scalable cloud platform with a Google-style search engine available across AbbVie's vast data.
- Integrate with CRM systems, data lakes, and data warehouses.
- Establish enhanced data governance and quality to boost user trust in data and simplify regulatory compliance.

Clearly, a transition of this size must be handled carefully. Recognizing this, AbbVie avoided a "big bang" rollout and chose to begin gradually. This careful

approach reduced risk and increased organizational confidence in the new platform we worked on together.

Our work in developing a modern data platform that integrated all their information proved to be timely. In 2020, when AbbVie acquired Allergan, another pharmaceutical company of similar size, it instantly doubled its scale. We know from previous experience and academic research that the expected benefits from mergers and acquisitions often do not materialize because of poor post-merger integration, and technology incompatibility is frequently the cause. It would have been the case with AbbVie if its multiple legacy platforms had still been in use. However, because we helped AbbVie build a strong data foundation and establish an integration hub, they were able to seamlessly incorporate Allergan's data remotely—during the pandemic, no less—with minimal disruption. Because of COVID-19, the entire integration was completed remotely, with no one in the AbbVie offices. In other words, this massive data merger was achieved entirely *remotely*.

AbbVie used Reltio's modern data unification platform to help make its post-merger integration with Allergan a success. Traditional mergers often struggle to create value because of technological incompatibilities and widespread data issues, resulting in operational delays and slower realization of synergies. In contrast, AbbVie's deployment of Reltio enabled a quick, effective data strategy that avoided the common pitfalls of older systems. The platform provided a scalable, cloud-native hub for data consumption and integration, allowing for rapid creation of a unified 360-degree view of customer, product, and transactional data across the merged organization. This efficient process ensured data quality and enhanced governance. It also accelerated the generation of actionable insights, reduced operational disruptions, and shortened time-to-value for the combined organization's commercial and clinical functions.

The stock market clearly has recognized the tangible benefits of this data-driven approach. Although initial reactions to the 2019 merger announcement

were skeptical, AbbVie's consistent execution and strong financial performance after the integration have since boosted investor confidence. The company's stock has outperformed the S&P 500 since the merger closed in May 2020, delivering returns of over 100 percent by December 2025.⁹ This positive valuation directly reflects the successful integration of Allergan's assets, made possible by a solid data foundation that enabled the quick realization of synergies. The financial market's appreciation highlights an important lesson: A modern, flexible data platform is a strategic necessity that directly influences an organization's ability to execute successful mergers and acquisitions. If you are considering acquiring a company to realize operational synergies, or if you believe you may be an attractive acquisition target, recognize that your data architecture could be a hidden advantage.

Adopt a Moonshot Mindset

In 1961, as the Soviet Union led the space race, President John F. Kennedy responded not with caution but with clarity. He declared that America would land a man on the moon and bring him back safely before the decade was over—a goal so bold it seemed impossible at the time. Yet that clarity of purpose became the guiding principle for everything that followed: budgets, scientific priorities, engineering decisions, supplier networks, and national morale. The moon landing was not just a technological achievement; it was the triumph of a vision that unified an entire system—and the nation—around a shared goal.

Agentic intelligence could be the moonshot for many legacy companies if—and only if—they treat data as a strategic asset rather than just an administrative byproduct. Take Tata Steel, one of the world's oldest steel producers. Confronted with global competition, fluctuating energy costs, and increasing sustainability demands, the company's leadership redefined data as the foundation for operational excellence—not merely an IT project but a core part of their corporate philosophy. Their strategy centered on creating a

“connected steel plant,” leveraging sensors, advanced analytics, digital twins, and AI-driven optimization to boost throughput, minimize downtime, and improve safety.¹⁰ Importantly, company leaders have positioned data as the new source of differentiation in a saturated industry. Their public messaging highlighted the goal of building a “digitally enabled, intelligent steelmaking ecosystem,”¹¹ where insights—not intuition—drive key decisions. This perspective elevated data to a board-level priority: Investment increased, cross-functional teams emerged, and talent development programs began training the next generation of digital industrial operators. Tata Steel isn’t just treating data as a tool; it views it as a moonshot for reinvention.

Walmart offers a different yet equally strong example. As the world’s largest retailer, Walmart realized early on that size alone wouldn’t guarantee future success. Its leadership deliberately shifted its strategy from focusing only on retail logistics to emphasizing retail intelligence. Under CEO Doug McMillon, the company invested significantly in integrating store, e-commerce, and supply chain data into a single system that provides real-time insights and supports quick decision-making. This strategic move led to the creation of advanced inventory systems, predictive replenishment models, in-store computer vision, and the integration of marketplace data into planning processes. Walmart’s leaders consistently emphasize that data and technology are essential, not optional, to their new operating model. The company aimed to build “the world’s most customer-centric, data-driven retail ecosystem.” This unified thousands of stores, hundreds of thousands of employees, and extensive supplier networks around a technological transformation that now positions Walmart not just as a retailer using technology for efficiency but as a leader in real-time retail intelligence.¹²

These two cases illustrate a simple but powerful truth for the Age of Agentic Intelligence: Strategy precedes architecture, and architecture precedes intelligence. Tata Steel and Walmart did not begin with advanced AI models. They started by recognizing that their data needed to function more like a

neural system—connected, contextual, and continuously refreshed. Once leaders committed to that vision, their organizations reorganized around it: Budgets shifted, governance modernized, silos weakened, and data strategy became business strategy. This is the leadership lesson of the moonshot: Bold clarity transforms capability. In the agentic era, companies that succeed will be those that commit first to an architecture that enables humans and AI agents to think, learn, and act together.

You may be aware of additional cases in which leaders have fully embraced the view that data is the foundation of agentic intelligence and are systematically developing strategies to transform themselves. The key point is that it's not just about adding AI models to legacy organizations; it's about rethinking organizational intelligence. We understand that human intelligence isn't solely dependent on the number of neurons in our brains—it's because all of them can *connect*. Signals fire across synapses at lightning speed, enabling patterns to form and actions to occur in fractions of a second. This allows us to think, reason, judge, and recall memories.

Now machines can imitate that phenomenon. Technology can now develop models that learn, adapt, and act as independent agents—making decisions, predicting results, and working alongside humans. In the age of agentic AI, algorithms do more than just crunch data; they *act* in the real world: advising doctors, managing inventories, negotiating contracts with suppliers, and detecting fraud, among other tasks.

But there is a catch. No matter how advanced the model, an AI agent is only as good as the data it can access. Remember, *intelligence depends on being able to connect the dots*. It's the only way AI agents can complete their tasks. But without the right architecture in place, these agents are like pilots flying with fogged-up windshields: They have powerful engines but are dangerously blind and unaware of nearby aircraft.

Any company that hasn't adopted the neural network approach to data will struggle to fully leverage the new, powerful AI tools. Businesses need to move

away from the traditional filing-cabinet mindset and develop a data architecture that allows every signal, from every part of the organization, to reach where it's needed instantly.

Clearly, building the right architecture is the foundation of data success and maximizing AI's potential. Data isn't just about dashboards or software; it's about connectivity. We need to move away from a filing-cabinet mindset and design data as a neural network—connected across the enterprise, constantly updating, and rich in context. That's how Ant Group's AI makes hundreds of millions of credit and risk decisions daily,¹³ how Netflix delivers highly accurate recommendations, and how Estonia operates a government in which citizens provide their data once, and it flows smoothly across services.¹⁴ (Think of that minor miracle the next time you visit a new doctor's office and fill out three pages of questions—using a pen on a clipboard—the same information you've shared in countless other medical visits.)

The lesson for leaders is that it's not about the size of your data lake but how smoothly and intelligently the data moves. A neural data architecture functions as the brain of the modern company—the backbone that enables both people and AI to sense, decide, and act in real time. Without it, even the smartest algorithms remain inflexible. With it, organizations develop resilience, agility, and the capacity to innovate at scale.


Our message is clear: *Those who organize their data like a brain will succeed.* Those who store it away will fall behind. The future of competitive advantage depends not on how much data you have, but on how smartly it flows. It all comes down to how well your company can *think*.

Think.¹⁵ That was probably the most famous one-word business slogan in history, coined in the mid-twentieth century by Thomas J. Watson, the legendary founder and CEO of IBM, whose surname is still used today for IBM's suite of AI services. For decades, IBM lived by that word because it not only resonated with its customers but also highlighted the computer giant's primary benefit to companies—the ability to use an electronic “brain” to

process data. Back then, it was a minimal brain, indeed, because it was based on the now-obsolete filing-cabinet mentality.

Now, it's an entirely new game. Yet, the importance of thinking persists. The potential of today's neural network architectures has brought us much closer to achieving the goal of replicating the human brain's functions—and scaling them for use by an entire corporation and its extended ecosystems.

OceanofPDF.com



RULE 2

DATA IN MOTION—NOT DATA AT REST—CREATES AGENTIC ADVANTAGE

Data is in our team's blood.

—CHRISTIAN HORNER, CEO OF ORACLE RED BULL RACING

Formula 1 (F1) racing captures our attention with its speed, and that's the exciting part. That's why Apple reportedly paid Brad Pitt around \$300 million to make a movie about it—and why the film became an international hit, earning over \$600 million at the box office. Additionally, Apple has agreed to pay \$150 million annually for US streaming rights to F1 racing.¹⁶

But, to us, what's more impressive about F1 racing is how quickly the *data* moves. Theoretically, data travels at the speed of light (about 186,000 miles per second), but it is often slowed by the medium it travels through (fiber, copper, or air), network delays, and other such factors. Typically, data travels at about two-thirds the speed of light (around 124,000 miles per second) in fiber-optic networks, with actual internet speeds determined by bandwidth, not just physical travel time. Want to discuss your car loaded with sensors and

software? There are three hundred sensors on each F1 car. During a race, 1.5 terabytes of data are collected to analyze and optimize up to four thousand parameters, generating one hundred thousand data points—all used to improve everything from car design to racing strategies. Because the data is transmitted in real time, it can support team decisions during a race, even as the car speeds past two hundred miles per hour. Instincts don't always work for drivers at that speed—but data does.¹⁷

That's because F1 is really a contest of who can close the loop fastest—sense, predict, act, repeat. AI services give drivers an incredible edge in that pursuit by analyzing car telemetry, pit-lane video, timing loops, and even radio chatter to detect anomalies, streamline decisions, and shave milliseconds—yes, milliseconds—off stops and out laps. Simultaneously, observability tools highlight performance issues, such as heat spikes, sensor drift, and clutch behavior, allowing engineers to intervene before a minor problem becomes a race-losing defect. Achieving the winning speed depends on orchestration—delivering the right signal to the right decision model and the right human (or agent) at the right time.

Yes, in F1 racing, speed is crucial. But now, the speed of that data is just as important as the car's speed.

F1 is not an outlier. Speed is quickly becoming a key differentiator in many industries worldwide. The business world has recognized that real-time data provides a strong competitive advantage, enabling companies to adapt quickly—precisely when the market demands a shift. *That's why data velocity, not data volume, now rules.*

Lowering Latency

Why is the shift to emphasizing velocity over volume so crucial? In one word: *latency*.

Latency is the delay between an action and its response—the time it takes for data, a signal, or a system process to move from one point to another. The higher the latency, the longer the delay. In the Intelligence Age, we aim for the lowest possible latency because low-latency systems bridge the gap between insight and action, transforming reactivity into proactivity. Managing latency across ingestion, processing, and delivery determines whether intelligence moves at market speed or lags.

In contrast, high latency means “hurry up and wait.” Queries take longer, visualizations and insights update more slowly on dashboards, and systems spend more time waiting for responses instead of processing new requests. Here’s where it gets dangerous: With high latency, *the data you’re seeing is often already outdated by the time it loads*. That undermines real-time monitoring, forecasting, and automation—and it turns what should be continuous intelligence into after-the-fact reporting.

From a business-user perspective, high latency feels like friction—waiting for files to load, systems to sync, or queries to return. This slows decision-making, discourages the adoption of data tools, and often causes people to revert to offline work-arounds. We observed this in our AbbVie case study, which we discussed in the previous chapter. Before AbbVie unified their data and implemented agentic intelligence, their latency was extremely high. Performance was sluggish, and users consistently reported data latency. And remember how long it used to take to input a record into all their legacy systems? *An hour*. That’s unacceptable today, and thankfully, they recognized that.

Can we quantify the actual cost of latency? Some studies show that: (1) every one hundred milliseconds of additional page load time results in roughly a 1 percent decrease in sales;¹⁸ (2) a one-second delay in page load time can cause a 7 percent drop in conversion rate, and a three-second delay can lead to a 20 percent decrease;¹⁹ and (3) workflow execution times increase by 100 percent to 550 percent when database or network latency rises.²⁰

Uber's reliance on minimal data latency is essential for its financial stability and operational success. Delays in data can directly affect the accuracy of vital functions, such as real-time estimated times of arrival (ETAs) and dynamic surge pricing calculations, which can lead to poor balancing of riders and drivers, decreased customer and driver satisfaction, and lost revenue opportunities. By aggressively reducing latency, Uber enhances infrastructure cost-efficiency by improving hardware utilization while ensuring a fast, reliable user experience that is crucial to customer retention and profitable on-demand services.²¹

A study funded by Google shows that reducing a key web metric by just one-tenth of a second can boost conversion rates by up to 8.4 percent in retail and 10.1 percent in travel and increase average mobile order value by 9.2 percent.²² That's a significant tenth of a second. Clearly, latency isn't just a technical concern—it *directly affects* business metrics such as revenue, customer experience, and operational efficiency, and offers a competitive edge.

Bigger is no longer better—speed is now the imperative. Each decade since 1990 has compressed time a little further—shrinking the lag between sensing change and acting on it. What used to take days now takes seconds or even milliseconds. What once required numerous planning meetings now happens automatically in the background through AI agents. Machines now monitor, analyze, and respond faster than human teams can convene to discuss what happened.

Today, the fastest organization wins—not just in data speed but also in its ability to learn, decide, and adapt smoothly. Time has become the new benchmark. First-mover advantage was the rallying cry of the Industrial Age. Now, it's fast-mover advantage—and this edge comes from how you organize data.

Why Data Latency Matters

The technical view of latency arises from three sources: (1) network—the delay from *sending a signal to receiving a response*, such as how long it takes for a web page to load after you click a link; (2) processing—the time a system takes to *analyze data and respond*, such as a trading algorithm reacting to market changes; and (3) storage—the duration needed to *access information* from memory or a database.

From a management perspective, there's another aspect to consider. It's *cognitive latency*: the delay in *interpreting and acting* on available data caused by human hesitation, role ambiguity, goal misalignment, or unclear accountability. And it warrants closer examination.²³ It raises an urgent and fundamental question: Can we keep up with the technology we've implemented?

While engineers can reduce network, processing, and storage delays, the real bottleneck often lies in human and organizational responses. Decisions slow down because managers question data quality, governance processes delay access, or teams hesitate to act without consensus. As more organizations adopt agentic systems, cognitive latency becomes the main constraint. Reducing it requires more than faster systems; it needs trustworthy governance, transparent data lineage, and clear decision rights so humans and intelligent agents can act confidently on the same information.

When all four types of latency are managed, we can operate at the same speed as our data, deliver exceptional customer experiences, and create meaningful engagement instantly. Whether in finance, manufacturing, healthcare, hospitality, or F1 racing, reducing latency can be the difference between winning and losing—because now, every millisecond matters. Today's consumers expect to get what they want with a tap. AI agents require the same level of efficiency in their queries. Companies that deliver this level of experience will have the best chance of thriving. You'll see that for yourself—because we're about to show you how data velocity has transformed several key

industries in ways that would have been unimaginable in the twentieth century.

Fast Fashion

Zara was among the first global companies to recognize that data speed matters more than data volume. Long before terms such as *real-time analytics* or *agentic intelligence* became common in business, Zara based its entire model on velocity. The company pioneered fast fashion not by out-designing competitors but by out-sensing and out-responding them. The key insight was that the apparel business wasn't about forecasting seasons months ahead; it was about spotting weak signals in consumer behavior and turning those signals into products *before* competitors even recognized the trend. Zara could move from sketch to store in two to three weeks—much faster than the industry average. This operational rhythm worked because Zara made speed of information its core strategy.²⁴

Crucially, Zara's strategic approach to data was fundamentally different from that of traditional retailers. Instead of relying on large historical datasets, the company empowered store managers to serve as real-time intelligence hubs. Every day, they shared observations with designers about what customers tried on but didn't buy, what they asked for but couldn't find, and which items were selling out quickly. Designers responded immediately—adjusting cuts, fabrics, and colors in real time. Zara's competitive edge wasn't just fashion intuition (which still prevails in many fashion houses); it was a continuously updated flow of high-velocity, high-context data. In today's terms, Zara created an analog version of a continuous feedback loop, where frontline signals directly influenced decision-making, and decisions quickly moved into production. Historical data mattered far less than what customers likely wanted *that afternoon*.²⁵

As consumer expectations shifted and sustainability concerns grew, Zara once again prioritized speed over volume in data use to stay adaptable. The company now employs AI and advanced analytics to enhance demand forecasting, minimize overproduction, and choose more sustainable materials for its “Join Life” collection.²⁶ These systems work not because Zara has more data than competitors—it doesn’t—but because it has mastered reducing the feedback loop between culture and commerce, insight and action. This is the deeper lesson. Modern organizations preparing for agentic intelligence often believe that gathering more data will give them an advantage. Zara proves the opposite: Competitive edge comes from *data that moves quickly*—systems that sense, decide, and adapt faster than rivals. In a world where agents will increasingly coordinate decisions across design, supply chain, and operations, the companies that succeed will resemble Zara—organizations where information flows fast, is continuously updated, and is closely linked to real behavior using real-time data. Zara’s success shows that the future belongs not to those who collect the most data but to those who act the quickest with the data they have.

Speed is the differentiator—not volume.

Fast Money

Fast data isn’t just fashionable; it’s profitable.

The financial services industry has also prioritized speed over volume—because it really had no choice. Mastercard processes approximately one billion transactions daily, with each settled within three hundred milliseconds.²⁷ In that brief window, companies aim to detect fraud before the process is complete; otherwise, they risk losing funds. Yes, historical data is important to the models, but clearly, the advantage lies in acting within that one-third of a second. Without that ability, banks and credit card companies risk losing billions of dollars in trust and revenue, simply because they can’t respond

quickly enough to prevent fraud. After all, closing the barn door after the horse has escaped isn't effective—and neither is flagging a transaction after the scammer has left a jewelry store and spent a thousand dollars using someone else's credit card.²⁸

On the trading floor, time is also money. At Goldman Sachs, everyone understands that milliseconds can make or break a fortune. That's the reason for its Marquee platform—to provide bankers and clients with real-time analytics instead of outdated figures. After all, why bother reacting to what happened at the previous session's closing bell? In today's markets, where exposures change by the minute, end-of-day reports are about as useful as yesterday's weather forecast. Institutional investors choose Marquee because it offers access to market data, analytics, and liquidity in a modern digital format. This, in turn, helps them act faster, make more informed decisions, and manage risk with greater transparency.

Asset managers now see velocity anew. Blackstone, the world's largest alternative asset manager, with over a trillion dollars in assets under management, once relied on quarterly surveys to value commercial property. Today, those three months can feel like an eternity. Now, it analyzes rental flows, transaction data, credit card spending, and even satellite imagery to update valuations in real time. The company knows that missing a shift by a month—let alone three—can cost billions or put assets at risk. While historical data remains useful, the speed of data signals is crucial. This rapid insight allows Blackstone to often spot early trends, such as rising inflation, by monitoring developments within its portfolio companies. Jon Gray, president and chief operating officer of Blackstone, put it this way: “Data science empowers us to make better investment decisions and drive greater value for our portfolio companies. It is an integral part of Blackstone.”²⁹

You remember the saying, “Time is money.” *Speed links time to money.*

Fast Cars

At the beginning of this chapter, we examined how data governs F1 racing. Since we're assuming none of you are driving over two hundred miles per hour on your daily commute, let's shift gears and look at how it affects the vehicles we all drive.

What's interesting to note is that real-time data has been a part of our driving experience from the get-go. After all, we've always been able to access specific data points in the moment, such as how much fuel the car has (gas or charger gauge) and how fast we're going (speedometer). You might view those dashboard perennials as early harbingers of the immediacy of today's data. However, as you can imagine, things have changed dramatically in this sector too—because a new generation of car makers and ride services has leveraged real-time data far beyond fuel and miles per hour, covering every conceivable area of vehicle use.

For example, data on vehicle safety and usage were, by necessity, mostly historical. Old-school automakers—think industry leaders that have dominated the scene for most of the twentieth century—kept detailed engineering archives, including decades of crash tests, safety reports, and laboratory data. This enabled them to identify design issues and improve safety retrospectively. However, this process clearly took time, during which drivers and passengers of potentially dangerous vehicles were unknowingly put at risk.

That timeline has now been drastically shortened by new-school carmakers such as Tesla, which enables each vehicle to act as a node, uploading sensor data and driver interventions.³⁰ Central models are retrained and sent back over the air (OTA). What one car learns, millions soon know. Unlike rivals that test in controlled environments and roll out updates through lengthy dealer cycles, Tesla's edge is data velocity: a feedback loop—and software updates—measured in weeks, not years.

That edge allows Tesla to collect a massive amount of data from its vehicles—arguably more than any other automaker. When it comes to safety, Tesla uses aggregated data to analyze crash events and near misses, improving both physical and digital safety features. OTA updates are often based on this analysis; for example, adjusting braking behavior or collision-avoidance algorithms.

Every Tesla on the road serves as a rolling data sensor, continuously collecting data from its sensors. This data helps Tesla train and refine its AI models for Autopilot and Full Self-Driving (FSD). For example, when a car encounters an unusual road situation (such as a temporary lane closure or a rare traffic sign), Tesla can flag it and use that footage to retrain its neural networks. The company calls this “fleet learning,” essentially crowdsourcing driving intelligence from billions—yes, seven billion and counting³¹—of real-world miles. Data also informs improvements to battery performance, range estimation, and thermal management.

And the benefits of Tesla’s systems don’t end there. Vehicle sensors continuously monitor the health of various components, such as the battery, motor, brakes, and software. This enables the carmaker to predict failures before they occur, allowing service centers to alert owners or proactively schedule maintenance. The company can also diagnose issues remotely and often resolves software-related problems without a physical visit. Tesla also analyzes real-world driving behavior and vehicle usage patterns to inform design improvements.³²

Then there’s auto insurance. Tesla’s distinct advantage in the insurance market comes from its vertically integrated approach, which leverages the extensive data collected directly from its vehicles. Unlike traditional insurers that depend on basic demographic info such as age, location, and credit history to assess risk, Tesla uses built-in sensors to monitor specific driving behaviors in real time. This data generates a transparent “Safety Score” for the driver,

which directly influences their monthly premium in the twelve states where the service is currently available.

This model creates a safety feedback loop, offering an objective, transparent way for drivers to reduce costs by actively improving habits such as maintaining a safe following distance and avoiding hard braking or aggressive turns. The company has analyzed more than thirty billion miles of fleet data to refine its algorithm, enabling it to accurately predict collision rates and potentially deliver significant savings for safe drivers compared with competitors. By providing a fully digital, integrated experience, Tesla aims to revolutionize how insurance is priced and delivered, moving toward a dynamic model that puts the driver directly in control of their premiums.

Tesla's recent success and the financial market's expectations for its future can be linked to its groundbreaking use of vehicle data to train AI, enhance safety, predict maintenance, design better products and new services (such as Robotaxis), and even transform insurance. Tesla's learning and improvement cycle—in which vehicle data is sent to the Tesla cloud, processed for neural network training, and then returned to the vehicle as OTA updates—sets it apart from traditional automakers.

New entrants to the automotive industry, such as Uber, DiDi in China, and Waymo (a unit of Alphabet), all focus on using real-time data to inform their strategies and operations. Their business models wouldn't be possible without the ability to control and manage data velocity. It's also reasonable to view vehicles such as Tesla's Robotaxis and Waymo's fleet as agents—though physical ones—that operate and navigate autonomously on real roads, transporting people and goods.

The speed of data is now just as important as the speed of a vehicle.

Fast Response

Sometimes, high-speed data can even save lives.

Well-designed neural networks can improve healthcare efficiency and responsiveness. Yes, hospitals may hold vast amounts of patient records, but what truly matters is whether the right data is accessible at the right bedside, in real time. During the COVID-19 pandemic, we saw a stark divide between those who could quickly update their data and those who struggled to unify siloed databases. For instance, some systems could only update hospitalization figures weekly; they were essentially fighting the pandemic in the dark. Meanwhile, others built dashboards that pulled admissions, lab results, and ventilator usage data hourly. They could pivot swiftly when needed, reallocating staff and equipment before crises peaked.

At Mayo Clinic, speed has become a quiet yet defining principle of excellence in care.³³ Even before the pandemic, the institution understood that vast archives of medical records were only useful if insights arrived before the moment of decision. A patient's life often depends not on the amount of data but on whether the right signal reaches the clinician at the right time. That's why Mayo has invested in systems that deliver information *directly into the flow of care*. Cardiac monitoring, for example, no longer requires patients to spend days in hospital beds hooked up to machines. Instead, wearable devices transmit high-resolution electrocardiograms remotely, and physicians review them in near-real time, often while the patient is at home. Diagnosis and treatment can begin within the critical window—measured in minutes, not days—helping prevent health deterioration and unnecessary hospitalizations.

The same focus on speed drives Mayo Clinic Platform,³⁴ which provides de-identified, long-term patient data to innovators under strict oversight. Start-ups and researchers can test AI models with real, clinically relevant data streams, shortening what used to take years of retrospective research into months of development. One project used these data feeds to predict multiple cancers years before traditional diagnosis. Another used streaming lab data to identify adverse drug reactions as they happened, not after they appeared in case reports. By emphasizing speed over volume, Mayo has created an

ecosystem in which clinicians, AI tools, and entrepreneurs respond to signals in real time, delivering faster value to patients.

Mayo's example illustrates that the power of data velocity extends beyond pandemics or crises. In everyday care, it can lead to faster interventions, shorter hospital stays, and earlier diagnoses. As AI agents become part of clinical practice, velocity underpins human–AI teamwork.

Lives are saved when data moves at the speed of care.

Fast Impacts

Today's consumers expect service within seconds and minutes, not hours and days.

That's why industries such as fashion, finance, motor vehicles, and healthcare are not the only ones being transformed by data velocity. The traditional insurance sector is also being renewed by data speed. Traditional underwriting often took weeks, with claims processed slowly and fraudsters exploiting those delays. Today, start-ups such as Lemonade are based on “insurtech”—using AI and big data to make the industry more efficient, data-driven, and customer-centered. At Lemonade, claims can be filed, verified against policies, cross-checked against fraud databases, and paid within minutes, resulting in satisfied customers and a substantial reduction in fraud. Winning in today's insurance market lies in real-time processes, not in filing cabinets full of past claims.

Even fast food has become faster. With Reltio's help, McDonald's has transformed how it uses data to offer more convenience and engagement, increasing customer satisfaction and loyalty through personalized offers, mobile ordering, and digital kiosks. Operational efficiency has also improved with AI-powered supply chain management, real-time performance monitoring, and streamlined restaurant workflows, which reduce waste and wait times. On a deeper level, data analytics has provided valuable insights into

customer behavior and market trends, enabling predictive modeling and better decisions for menu development and marketing. Cloud computing, AI, machine learning, and big data analytics all support McDonald's agility, scalability, and ongoing innovation, helping the company respond quickly to changing customer needs and market demands.

What enables this level of data velocity? In the previous chapter, we introduced the concept of datagraphs, which connect entities—such as customers, products, contracts, and events—into a dynamic web. At Reltio, our datagraphs have enabled a millisecond-focused approach across various areas, providing our customers with a single data solution that integrates long-term and short-term memory; real-time, interactive search using natural language grounded in conversational experience; and tools to interact with and manage the data.

You might ask: “Haven't we always prioritized data velocity but been limited by the available technologies at different points in history?” In a way, yes. But AI agents add an extra dimension—their speed isn't constrained by human boundaries but by machine capabilities. They automate analysis, filter out noise, and make real-time decisions by processing and interpreting data instantly—whether it's sensor data, transaction logs, or customer interactions. They can also distinguish relevant signals from noise by recognizing patterns and prioritizing what's most important. This helps executives and systems focus on meaningful changes rather than raw data that offers little insight, and keeps attention and resources sharply aligned with what's time-sensitive or high-impact.

AI agents also continually learn from new data, refining predictions, detecting patterns earlier, and responding more quickly over time. This creates a feedback loop in which speed drives learning—and learning, in turn, drives more speed. In addition, they act as connectors between siloed data systems. They can translate, synchronize, and route data between departments or

platforms in real time, so there's no waiting for manual handoffs or overnight data refreshes.

Finally, AI agents not only move data faster but also make it *smarter* as they do. They transform raw information into real-time intelligence, enabling organizations to operate at the speed of change. This keeps AI functioning at the same pace as the real world. For leaders, the key point is clear: Data volume is meaningless without movement. Flow is what keeps a business alive, and agentic intelligence is what drives that flow. With the right data architecture, AI becomes incredibly powerful. Without it, AI often struggles even with the simplest tasks.

The Speed of Culture Change

We could share many more examples of how data velocity transformed traditional companies and helped start-ups succeed. One key factor that sets apart companies that see speed as an advantage is organizational culture. Others are stuck in established processes and ways of working that have served them well for decades. They do not see the need for change; they view AI agents as just another tech shift. Just as Kodak couldn't let go of physical film and Blockbuster couldn't abandon physical media, many businesses struggle to move away from their traditional methods—simply because they are unsure of the real benefits.

We tend to value what we already know more than what we don't, and old-fashioned recordkeeping systems are much more familiar to managers than datagraphs. Changing a company's fundamentals can evoke genuine fear because it introduces uncertainty, which is perceived as risky, especially when the current approach has *been effective*. Psychologists call this the status quo bias: the preference to maintain the status quo, even when better options are available; Venkat called this the “competency trap” in his 2023 book, *The Digital Matrix*.³⁵

We urge leaders not to be stuck in their traditional view of data speed. Improving data velocity provides many better options for any business.

Consider how most organizations previously operated. Decisions moved slowly upward through layers of hierarchy, with managers hoarding information as a source of power. Reports were gathered, sanitized, and delayed, so by the time executives received them, the situation on the ground might already have changed—and the window for action might have closed. However, a CEO might not have realized this because they were working with outdated information instead of the most current data.

In contrast, companies fully diving into the Intelligence Age are building digital cockpits—real-time dashboards and command centers that give everyone from the C-suite to frontline managers access to the same signals simultaneously. At Procter & Gamble, for example, “Business Spheres” display global performance metrics on wall-sized screens, enabling cross-functional teams to resolve issues in hours rather than weeks. Instead of debating whose spreadsheet is correct, leaders start from a shared truth and focus on action. Conflicts are diminished simply because everyone’s working from the same data. This becomes the necessary foundation for agentic workflows.

Velocity also changes how companies connect with the outside world. Firms that once depended on quarterly surveys can now monitor social media sentiment in real time. Listening centers at brands such as Dell and Starbucks gather customer feedback instantly, alerting teams right away when a product is trending, a complaint goes viral, or a crisis is emerging. Decisions that used to go through marketing and PR channels can now be made immediately, often with AI agents detecting patterns and suggesting responses. The cultural impact is significant: Teams feel empowered to act without waiting for approval, and executives develop a natural ability to adapt quickly.

Command centers are applying this approach to critical operations. For example, at airlines, integrated control rooms combine weather data, aircraft telemetry, crew schedules, and passenger flows into one view. Everyone—from

pilots to gate agents—works from the same real-time picture. Instead of silos competing over partial information, teams coordinate like neurons firing in unison.

The listening and command centers of the past have evolved from impressive physical rooms of the sort you'd see in a James Bond movie into virtual, distributed ecosystems driven by data velocity. Instead of a single war room with analysts viewing a data wall, AI-powered software now processes vast, high-velocity data from diverse digital channels, delivering immediate, contextual insights directly to teams. This evolution transforms the command center from a fixed location into a ubiquitous, collaborative capability, enabling agile, informed decisions anywhere, at any time.

Our message is this: You can't step up to compete in the Intelligence Age without focusing on data velocity. It requires leaders to analyze the root causes of high latency across the four categories we identified earlier. And of all those four categories, remember that it's cognitive latency that's most likely to hold companies back—particularly when humans and AI agents work collaboratively. While F1 drivers are intrinsically tuned to minimize cognitive latency (otherwise, they wouldn't be competing at such a level of intensity), most organizations are saddled with high cognitive latency.

So yes, fast data can open many surprising doors. In an era when AI agents work with humans in real time, slow data is worse than no data. Gaining a competitive edge depends on minimizing the gap between signal and decision. Speed sets the new standard. Data volume without speed is essentially useless.

The Industrial Age rewarded size. The Intelligence Age rewards speed. Just ask any F1 driver.



RULE 3

AGENTIC INTELLIGENCE REQUIRES CROSS- FUNCTIONAL, CROSS-PARTNER DATA

We must, indeed, all hang together or, most assuredly, we shall all hang separately.

—BENJAMIN FRANKLIN

It's checkup time.

You go in for your annual physical, and your doctor orders a few blood tests—routine stuff. A week later, he refers you to a specialist.

You visit the specialist's office, where you need to fill out pages of paperwork that include all the information you've already provided to your primary care physician and healthcare system. When you finally see the specialist, she asks many of the same questions you've just answered because she hadn't had time to review them. Then she finally gets to the point. After a quick evaluation, she recommends you get some blood work done.

You blink in confusion. “I just had that done.”

She checks her screen and says, “Hmm. I don’t see it here.”

So, a few days later, you’re back at the lab, rolling up your sleeve to get jabbed for the same blood work. OK. No big deal. But your insurance company denies coverage for the second test because, according to their records, it wasn’t authorized. You call your primary care doctor’s office. They say they sent the paperwork; it’s not their fault. Then it turns out they faxed it—to the wrong department.

Faxed? They still use *faxes*? you wonder, amusedly—just for a moment.

By the time it’s finally sorted out, you realize how much time you lost dealing with all this nonsense. You had to spend hours on hold before you could talk to someone. You had to answer the same questions about your medical history multiple times. You had to get blood work done twice. And the second time, you were expected to pay for it. All because data systems *don’t communicate with each other*.

And you’re the one stuck with the soul-crushing job of connecting the data dots for them.

Lessons from AbbVie

AbbVie offers a powerful illustration of what it takes to overcome one of the most persistent barriers to agentic intelligence: fragmented data within the enterprise and across its broader ecosystem. From the outset, the company recognized that meaningful progress in drug discovery would be impossible as long as data remained trapped in scientific, organizational, and institutional silos. Its response was not a single platform or technology fix but a deliberate, multilayered integration strategy designed to make data seamless, trusted, and actionable.

Internally, it created the AbbVie R&D Convergence Hub (ARCH), moving beyond isolated data warehouses to a unified knowledge platform.

ARCH integrates data from more than two hundred sources, spanning early discovery chemistry, translational research, clinical trials, and manufacturing. As a result, a scientist working in late-stage clinical development can directly access raw experimental data generated years earlier in discovery, enriched with knowledge graphs curated by subject-matter experts. What once required manual reconciliation across discovery, clinical operations, and manufacturing is now available within a shared context, enabling cross-functional collaboration and advanced, AI-driven analysis.

Externally, AbbVie recognized that understanding disease progression and patient outcomes requires data beyond its own laboratories. The company, therefore, built an extensive partnership network—supported by dedicated search, business development, and alliance teams—integrating data from more than 250 active partners under rigorous governance and compliance standards, including EU model contracts. A cornerstone of this effort is the AbbVie Research Collaborative, which enables patients to voluntarily contribute real-world data such as genomic information, electronic medical records, and lifestyle data. These insights add critical context that traditional clinical trials often miss. AbbVie also participates in reciprocal data-sharing initiatives, contributing anonymized datasets while integrating public and academic research through platforms such as Vivli and partnerships with organizations such as Calico Life Sciences.

The result of these internal and external efforts is an integrated data fabric that connects laboratory insights, clinical outcomes, patient experiences, and external research into a coherent whole. This foundation is what makes AI truly effective at AbbVie. ARCH does not simply store data; it provides the structure, provenance, and context that AI agents require to operate reliably. By embedding datagraphs into the platform, AI systems can interpret information within its scientific and clinical context rather than treating it as undifferentiated input, making their outputs more trustworthy and actionable.

This environment delivers three tangible benefits. First, AI acts as a force multiplier: Human experts no longer spend time reconciling disparate datasets, instead working alongside AI agents that synthesize insights across internal research, external science, and real-world patient data. Second, the platform enables rapid experimentation, allowing data scientists to deploy and test new models and specialized agents; for example, to identify novel drug targets or optimize chemical synthesis pathways. Third, it enables true human–AI partnership: An AbbVie scientist can engage directly with an AI agent that proposes hypotheses drawn from millions of data points, while applying human judgment to validate, refine, or challenge those insights.

AbbVie’s journey underscores a broader lesson: Connecting to external data ecosystems is difficult, especially in regulated environments where accuracy, privacy, and trust are paramount. But isolation is no longer a viable strategy. AbbVie showed that standards do not need to be lowered to enable sharing; with the right governance, architecture, and mindset, organizations can protect trust while accelerating learning.

In the Intelligence Age, connection—not isolation—is the source of advantage. Companies that invest in integrated data foundations across internal and external boundaries are better positioned to deploy agents effectively, respond to change faster, and translate intelligence into impact. AbbVie’s experience demonstrates that agentic intelligence is not unlocked by AI alone, but by redesigning how data, humans, and machines work together at scale.

To grow your business, start breaking down the data barriers.

Stepping Up to Open the “Data Box”

For decades, companies have treated data like a sealed box—valuable, private, and tightly protected. That instinct made sense. What was inside the box—customer histories, operational insights, proprietary knowledge—represented

control and competitive advantage. In an industrial economy, owning data was power. But in the Intelligence Age, that logic is becoming a liability. Value no longer comes primarily from hoarding information but from connecting it—across functions, partners, and ecosystems—to enable learning and action in real time. Holding on to the sealed box today is less like safeguarding an asset and more like dragging an anchor.

The companies moving fastest are not just opening the box. They are doing something more deliberate: opening it just enough to allow data to move, combine, and learn—securely, selectively, and with purpose. The objective is not openness for its own sake but smart connection to gain an edge. This is where many leaders hesitate, and understandably so. Opening data raises real concerns about risk, standards, and trust. Who sets the rules? How is misuse prevented? How do you collaborate without losing control?

The answer is not blind sharing but engineered trust. The most effective organizations start small, anchoring data sharing around a clearly defined, high-value use case where incentives are aligned. They invest early in governance, not as a compliance overlay, but as a core design requirement. Trust in these systems is not assumed; it is built into the architecture. When trust is automated, auditable, and enforceable, it scales.

This is where mechanisms such as data passports, federated learning, and shared governance frameworks matter—not as technologies, but as enablers of coordination. A data passport, for example, ensures that data carries its provenance, permissions, and usage history wherever it flows, making collaboration possible without ambiguity. Frameworks such as India's Account Aggregator system show that large-scale, secure data exchange can work when consent, traceability, and accountability are engineered from the start, enabling banks and fintech firms to coordinate around customers with precision rather than suspicion.

The strategic shift is subtle but profound. Competition moves from owning more data to activating trustworthy data across ecosystems. When data

can move with verified trust, new forms of coordination become possible: financial services that combine identity and transaction context in real time, healthcare systems that integrate clinical and patient-reported data responsibly, or mobility platforms that coordinate vehicles, insurance, and service histories dynamically. These are early expressions of agentic coordination—multiple actors, human and machine, acting on shared context.

Three forces make this shift unavoidable. First, modern AI systems learn best from diverse, connected data; intelligence trained in isolation is inevitably narrow. Second, the most important problems companies face—resilience, risk, sustainability, personalization—are systemic and cross-organizational. No firm can solve them alone. Third, advances in privacy-preserving technologies now allow collaboration without surrendering control, making isolation a strategic choice rather than a necessity.

In the Intelligence Age, connection is not the enemy of advantage; it is its source. Companies that learn how to participate in trusted data ecosystems—without compromising standards—will be better positioned to build agentic services that sense context, anticipate change, and act across boundaries. Those that cling to the sealed box may feel safe, but they will move more slowly in a world that no longer waits.

Data in Motion Is the Heartbeat of Agentic Agriculture

Digital transformation in agriculture is not new. For decades, governments, industry groups, and agribusinesses have promoted “farm-to-fork” and “field-to-market” initiatives, hoping data would flow smoothly from soil to supply chain to supermarket shelf. Yet despite years of effort, these initiatives delivered limited gains in productivity or sustainability. The problem was never a lack of data. Farmers have always generated rich information—soil maps, weather records, seed prescriptions, machine logs, yield monitors. The problem was

that this data sat in silos: locked in proprietary systems, fragmented across partners, constrained by poor connectivity, and analyzed long after decisions had been made. Without data moving in real time across machines, fields, and markets, insight arrived too late to matter.

That is now beginning to change. A new generation of agricultural companies—John Deere, CNH Industrial, and AGCO among them—are redesigning farming around *data in motion*, not data at rest. Their platforms share a common logic: Instrument every machine, connect every operation, and combine agronomic insight with equipment intelligence. This is not simply digitizing agriculture; it is enabling coordinated, real-time action. Planters adjust seed density meter by meter based on soil conditions. Sprayers adapt droplet size in response to weather and canopy conditions. Combines alter threshing parameters continuously as field conditions evolve. These decisions are made not by isolated machines but by connected agents—software and hardware working together on live data streams.

The critical shift is from intermittent to continuous connectivity. Advances in satellite networks, cellular IoT, and edge computing now allow machines to send and receive data in real time, even in remote fields where broadband was once unavailable. Farming has become a real-time optimization problem rather than a seasonal one. Instead of analyzing outcomes after harvest, farmers and their digital agents adjust actions as conditions change. Equipment issues are flagged before downtime occurs. Input levels are corrected instantly. Plans adapt automatically when weather patterns shift. This is what data velocity enables: the ability to sense, decide, and act while it still matters.

Real-time data also unlocks something that earlier farm-to-fork initiatives could not: *ecosystem-level intelligence*. Agronomists, seed companies, processors, insurers, and equipment dealers can participate in governed data-sharing environments where farmers control permissions and lineage. Rather than exchanging static files at season's end, partners access trusted, high-frequency data streams. Insurers can price risk based on real-time indicators rather than

historical averages. Food companies such as Unilever and Nestlé can verify sustainability claims with granular, auditable evidence. Supply chains that once relied on estimates can observe progress directly.

This is the emergence of *agentic agriculture*: a system in which humans, machines, and algorithms collaborate continuously around shared context. Tractors and sprayers become digital agents acting on the farmer's behalf. Decision-support tools evolve into real-time advisors. Sustainability shifts from an aspiration to a measurable practice, as the system can finally observe and respond to environmental change.

Why does this matter? Because the future of food demands more than incremental improvement. Agriculture must produce more with fewer inputs, adapt rapidly to climate volatility, and demonstrate environmental stewardship with confidence. None of this is possible with data locked in silos or analyzed weeks too late. It requires data in motion, supported by trust and interoperability across the ecosystem.

The failure of earlier farm-to-fork initiatives was not due to a lack of ambition but to a lack of data flow. Agentic agriculture is only achievable if data flows quickly, reliably, and seamlessly.

What the “My Data” Versus “Our Data” Fight Taught Us

In early 2020, as the world went into lockdown, the global auto industry reacted as it always had in downturns. Showrooms went dark, demand collapsed, and automakers from Detroit to Tokyo sharply cut semiconductor orders. Each firm relied on its own forecasts, guarded its production plans, and sent muted signals through layered supplier relationships. No one saw the full picture—and no one expected demand to return quickly.

But it did. By late summer, consumers confined to their homes began buying vehicles again, especially SUVs and pickups. At the same time, demand

for laptops, gaming consoles, networking equipment, and smartphones surged. Semiconductor manufacturers saw clear, continuous signals from consumer electronics firms: orders flowing in real time, backed by visible demand. From the auto sector, by contrast, signals were weak or absent. With limited visibility into automakers' rebound, chipmakers rationally reallocated capacity to customers showing immediate, verifiable need.

By the time automakers reentered the market, capacity was gone—and the crisis was already underway.

Across the industry, nearly finished vehicles accumulated in vast parking lots: engines installed, paint complete, interiors finished—missing only a few low-cost chips controlling critical modules. Ford suspended production. General Motors cut SUV output. Toyota, long celebrated for just-in-time excellence, announced sweeping production reductions. Volkswagen warned investors of lost volume. Nissan stripped features from vehicles to keep lines moving. What crippled one of the world's most advanced manufacturing systems was not a shortage of steel or labor, but the absence of a handful of semiconductors costing only a few dollars.

The economic impact was staggering. Dealership inventories vanished. Used car prices soared. Consumers waited months for delivery. Analysts later estimated that the industry lost roughly eleven million vehicles and nearly \$210 billion in revenue in 2021 alone. Technology had not failed the industry; coordination had.

Postmortems revealed a deeper problem. Automakers lacked real-time visibility into their own demand recovery. Tier 1 suppliers shared forecasts selectively. Semiconductor manufacturers had no integrated view of production plans across original equipment manufacturers (OEMs). Each firm optimized locally with incomplete information, while the system failed. Forecasts existed, but shared intelligence did not. The industry had built a highly complex global supply network on partial truths and delayed signals.

Executives eventually acknowledged what had long gone unsaid: The crisis was not simply a chip shortage—it was a shortage of shared context. Data had been treated as a source of bargaining power rather than as a means of coordination. Secrecy was mistaken for strength. In reality, it created fragility.

By 2021, that mindset began to shift. Automakers formed cross-functional war rooms to aggregate demand signals. Suppliers created secure channels to share projections more frequently. Chipmakers pushed for continuous visibility into needs and constraints. Quarterly spreadsheets gave way to digital interfaces designed to synchronize plans as conditions evolved. The goal was no longer to protect information but to align it.

The lesson extends far beyond semiconductors. In the Intelligence Age, AI agents cannot coordinate across boundaries without shared context. Autonomous planning systems fail when fed only internal data. System-level resilience cannot emerge from company-level secrecy. What the auto industry learned—at extraordinary cost—is that *“my data” cannot protect a tightly connected system; only “our data,” shared responsibly and in real time, can.*

Data hoarding, once a competitive instinct, has become a strategic vulnerability. Advantage now comes from architectures that allow organizations to see the same signals, at the same time, with shared meaning. That is the foundation on which agentic intelligence—and resilient ecosystems—are built.

Leaders: Open the Data Box Wisely

Sharing data across teams, partners, or even industries does not mean inviting chaos. Interconnection is risky when it’s done without structure. When organizations share data responsibly, three quiet foundations must line up behind the scenes: *meaning, rights, and trust*. These invisible layers keep complexity from becoming confusion.

Meaning begins with a shared language. Humans need shared words; machines need shared semantics. Without that foundation, data from one

system becomes gibberish to another. That's why industries build common standards—such as the Asset Administration Shell for industrial twins or the Proof of Activity Compliance Technology Protocol for carbon footprints. These standards act like dictionaries. They allow a supplier in India to report emissions in a format that a buyer in Germany can understand instantly. Leaders who grasp this don't start by asking, How do we share data? They start by asking, How do we ensure we're speaking the same language? Tata Steel, for example, spent years harmonizing the semantic meaning of its operational data across plants before inviting technology partners to participate. The result is a far more resilient and transparent supply network—not because they share more data, but because they share understandable data.

Rights come next. Think of them as digital treaties. Data contracts define who can use which data, under what circumstances, and for how long. In advanced ecosystems—such as the European data spaces—these rights travel with the data itself. Permissions aren't buried in PDFs or lost in inboxes; they are encoded as machine-readable rules. When Siemens opened its industrial cloud platform, it didn't ask partners for blind trust. It created reusable data contracts that made rights explicit and enforceable. In doing so, Siemens eased one of the greatest anxieties in any ecosystem: the fear that collaboration will become exploitation.

And then there is *trust*—not the soft, interpersonal kind, but institutional trust built through governance. Good governance is not a committee, a dashboard, or a compliance memo; it's a new logic of working. It defines standards, tests conformance, enforces accountability, and provides independent oversight. When Walmart designed its food traceability system with IBM, it wasn't trying to collect more data; it was trying to build a trustworthy network where every participant—from mango farms in Mexico to processors in the US—could contribute safely. The governance model made that possible. Trust emerged not from optimism but from audit trails, shared protocols, and transparent checks on power.

When these three layers—meaning, rights, and trust—are in place, sharing data doesn’t erode power; it amplifies intelligence. A “data commons” is not about giving your information away; it’s about allowing your data to *learn* by connecting with others. Industrial-age companies were built on silos and pipelines. Intelligence-age companies will be built on graphs and commons—structures where information flows, connects, and enriches itself through relationships.

So, the strategic question is no longer “Should we share data?” but “How do we design the conditions that make sharing safe, fair, and valuable?” How do we create the grammar, the contracts, and the governance that allow intelligence to move across boundaries without losing its integrity?

The leaders who take this challenge seriously are already showing what’s possible. In healthcare, Mayo Clinic’s platform strategy allows vetted partners to train AI on clinical insights without ever receiving identifiable patient data—a model that blends openness with rigorous governance. In energy, Ørsted shares turbine performance data with selected partners to accelerate predictive maintenance across the sector, knowing the gains come not from secrecy but from shared learning. In finance, India’s Account Aggregator framework—once opposed by banking CEOs—now powers a flourishing ecosystem of credit innovation by giving customers control over their data flows.

The organizations that master this balance—structured openness, disciplined interconnection—will define the next decade. They will shift from efficiency to adaptability, from isolated analytics to shared understanding, from digital transformation to collective intelligence. Opening the data box is only the first act. Designing how it connects—that is leadership.

* * * *

As we conclude this chapter, we want you to examine a familiar artifact—*an airline ticket*. Most of us think of an airline ticket as nothing more than proof that we have a seat on a flight. But the modern ticket is far more than that. It is

a doorway into one of the most sophisticated, interconnected data ecosystems on the planet—a quiet miracle of coordination that we barely notice anymore.

The moment you book an airline ticket, dozens of independent systems start working together. Reservation platforms update seat maps. Payment networks approve transactions. Loyalty programs track miles. Baggage systems generate tracking IDs. Border agencies screen passenger data. And once the aircraft pushes back from the gate, live weather data, aircraft telemetry, predictive maintenance systems, and air traffic control all come together to keep the flight moving safely and smoothly.

No single airline controls all of this. What makes it possible are the shared data standards developed over decades by the International Air Transport Association, the International Civil Aviation Organization, airport authorities, global distribution systems such as Sabre and Amadeus, and alliances such as oneworld, Star Alliance, and SkyTeam. These shared grammars and protocols turn thousands of independent organizations into a single functioning network. Without them, every step of travel—booking, transfers, baggage, boarding—would be a tangled mess of incompatible systems and manual work-arounds.

For leaders, the airline industry is not just a lesson in interoperability; it's a preview of the agentic future. Once data starts to flow reliably across organizations, human and digital agents can collaborate on behalf of the traveler. Imagine a traveler's personal AI agent coordinating with the airline's rebooking agent, the airport's congestion agent, the hotel's check-in agent, and even a city's ground transportation agent. Flight delayed? Your agent automatically adjusts your connection, notifies your hotel, and secures a later ride. Bag rerouted? Your agent tracks it in real time and updates you before you even land. Weather disruption? Agents collaborate across airlines to re-sequence flights, gates, and crews before chaos sets in.

This is not science fiction. It is the logical next step once data interoperability becomes a real experience. Travel already consists of an

interconnected system—agents simply make that coordination visible, proactive, and personalized. They cut down inefficiencies by communicating instantly across entities that currently rely on phone calls, emails, and hope. They also improve service levels by predicting problems and fixing them before travelers ever notice the friction.

The airline ticket, therefore, embodies the full arc of this chapter. It begins with shared data. It evolves into shared meaning. It deepens through trusted governance. And it culminates in *agentic coordination*, where human and digital agents work together to deliver seamless journeys despite the hundreds of things that can (and often do) go wrong.

Airlines learned long ago that no organization can deliver end-to-end value alone. They built the foundational interoperability that allowed intelligence to flow. Now, as industries everywhere prepare for the Intelligence Age, the same lesson applies: *Interoperability is the runway; agentic coordination is the flight.*

OceanofPDF.com



RULE 4

TRUSTED DATA IS THE NEW CURRENCY OF COAGENCY

We need to look at the outputs of gen AI with a critical eye and apply our human judgment to evaluate whether we trust them ... As long as we keep humans in the loop, we can do a lot of good with gen AI.

—BEN ELLENCWEIG, MCKINSEY

When it comes to data and agentic intelligence, trust is all-important.

Just after sunrise on a clear January morning in 2023, a Tesla Model Y glided south along Highway 101. The traffic was light enough that the driver let the automated system handle most of the driving. It was the kind of calm, ordinary moment that made the future seem almost routine—an intelligent machine quietly sharing the task of getting someone where they needed to go.

And then, without warning, the car slammed on its brakes.

There was nothing ahead. No stalled vehicle. No fallen tree. No hazard the driver could see. Yet the machine reacted as if an object had appeared in its path, and traffic behind it lurched into chaos as cars swerved and collided. Later, investigators called it “phantom braking,” a term now familiar to anyone

working with autonomous systems. The car took an action that made perfect sense *based on the data it received*—but not to the human inside.

When engineers examined the logs, the story became clearer. A moving shadow on the road had been briefly misclassified as a solid obstacle. The system followed its logic perfectly. However, that logic was based on a flawed view of reality. Waymo has reported similar hesitation events in its early public-safety analyses, where brief sensor noise or conflicting signals caused unexpected actions. The main issue was never “AI hallucination,” as headlines sometimes suggest. It was something more technical and more human: *The machine was operating on a version of reality that couldn’t be trusted.*

Anyone working in autonomy understands this inherently. As auto leaders often highlight, the toughest part of autonomous driving isn’t the sophistication of the algorithms; it’s the *quality and consistency of the data that feeds them*. A single misclassified frame, a corrupted lidar return, or a poorly timestamped sensor measurement can trigger a cascading effect throughout the perception and decision-making system. One unreliable data point can disrupt the collaboration between humans and machines.

And once that trust fractures, the partnership—and the *coagency*—collapses.

This incident is not simply a cautionary tale about self-driving cars. It is a mirror held up to the entire Intelligence Age. In every industry where humans and machines work together—as coagents—trust in data becomes the shared foundation of their relationship. When the data falters, so does the relationship.

Change the setting, and the stakes shift, but the pattern is constant.

In a hospital, an early warning system meant to detect sepsis can trigger false alarms—or miss real ones—when vital-sign feeds are delayed, incomplete, or misformatted. Nurses quickly learn to override the system. Doctors lose confidence in it. The benefit of coagency disappears not because the algorithm is weak, but because the *data pipeline* cannot be trusted to reflect reality. We

are sure you know this, as there have been occasions when you felt you couldn't fully trust the information or suggestions from the system.

In global aviation, advanced crew-scheduling systems aim to coordinate pilots, flight attendants, aircraft, and routes. When a major US airline faces a winter storm disruption, outdated availability data and inconsistent operational feeds can cause cascading errors. Flights that were already canceled might show up in rotation; crews stranded in other cities could be marked as available. Ground agents and cloud-based algorithms no longer interpret the same information, eroding trust among humans and in the system's authority. The outcome is paralysis.

Focus on finance, and the situation becomes even more detailed. Fraud-detection agents rely on real-time transaction logs. If timestamps shift, if a set of records duplicates, if a merchant's data feed breaks—even momentarily—the system starts seeing false threats everywhere. Humans can intervene, but suspicion grows: Which alerts are real, and which are false signals caused by faulty data? Coagency becomes adversarial when the shared understanding of truth breaks down.

Across all these domains, the same principle emerges—*humans and machines can only act together if they are standing on the same trustworthy data.*

Trust is not a layer on top of intelligence. Trust *is* the operational foundation—the invisible contract that makes coagency viable: “We see the world the same way; we act from the same source of truth.” The organization must act with a shared understanding of the context.

The autonomous car that braked for a ghost serves as a stark, cinematic reminder that even advanced AI—quick, accurate, sophisticated—can become dangerous if the data it relies on shifts, gets corrupted, or falls out of sync. But the bigger lesson is this: As more organizations deploy agents to coordinate supply chains, manage hospital floors, monitor financial risks, or support customer interactions, the question is no longer, How intelligent are the algorithms? It is, How trustworthy is the shared context they rely on?

Coagency is not about replacing humans or ceding authority to machines. It is about creating a partnership—fluid, reliable, mutually reinforcing. And that partnership begins at the most foundational level: the integrity, lineage, timeliness, and consistency of the data flowing through it.

The story of the phantom braking event is powerful not because of its drama but because of its universality. Every industry, every organization, every leader preparing for the agentic future must confront this question: *Do your systems—human and machine—see the same truth at the same time?* If the answer is no, then coagency isn't possible. If the answer is yes, then you have built the foundation for the only kind of intelligence that will matter in the decade ahead: *trustworthy, coordinated, human-machine intelligence grounded in shared reality.*

A decade ago, *data trust* meant worrying about breaches or leaks. Today, it has a deeper meaning: *Can humans and machines collaborate on a shared, accurate, and reliable view of the world?*

The autonomous car that braked for a ghost underscores that even advanced AI can be dangerous if it relies on unreliable data. As intelligent agents begin to assist in hospitals, power grids, logistics networks, and daily life, the importance of trust in data—not just algorithms—increases significantly. That's why trust in data will determine whether coagency becomes a competitive edge or a systemic risk.

Coagency Is Core to Agentic Intelligence

We must clarify what we mean by coagency and why it matters to the broader transformation toward agentic intelligence in enterprises.

Agentic intelligence is not simply using AI or overlaying AI intelligence on humans working the old way. It is an organization's ability to act intelligently through the coordinated efforts of human and machine agents.³⁶ In traditional automation, machines replaced human judgment. In agentic intelligence,

machines *enhance* human judgment, and humans *direct* machine actions. The relationship is reciprocal, dynamic, symbiotic, and ongoing.

But this reciprocity has a precondition: coagency. Humans and machines must be able to operate as partners who share context, coordinate actions, and maintain mutual understanding. An agent—whether human or machine—is an entity that perceives its environment and acts upon it with some degree of autonomy. Coagency emerges when multiple agents align their perceptions and actions toward shared goals.

The key insight we teach and practice with executives is that *trust in data underpins coagency, the mechanism by which agentic intelligence is realized*.

Without trustworthy data, you cannot have coagency. Without coagency, you cannot have true agentic intelligence—only isolated pockets of automation that humans work around rather than with. The enterprise remains stuck in what we might call “asymmetric agency,” where humans compensate for unreliable machines or machines operate in isolation from human insight.

Consider healthcare. Many hospitals now use predictive systems to alert clinicians to dangerous conditions such as sepsis. We can imagine what happens when an early warning algorithm is introduced into a hospital system with high hopes—until nurses start receiving alerts at the wrong times or not at all. The issue might not be the model’s math; it could be that the heart-rate and oxygen-saturation data arrived out of sync. Vital signs might have been delayed or overwritten when a patient moved floors or switched monitors. The result? Nurses lose trust and stop using the system; the physicians follow suit. The system might continue to operate, but collaboration becomes ineffective when the data can no longer be trusted.

Let’s look at what happened during a major winter storm in 2022 when Southwest Airlines’ scheduling system failed catastrophically.³⁷ Pilots stranded in one city appeared as available for flights elsewhere in the system. Flight attendants who had checked out for the day were listed as ready for duty. Aircraft that had been canceled hours earlier still appeared in the rotation. The

airline canceled more than sixteen thousand flights over ten days, leaving two million passengers stranded. The main problem was Southwest's outdated crew-scheduling software, which couldn't keep up with the chaos in real time, leaving pilots and crew out of position. This technical failure led to thousands of flight cancellations, overwhelmed the system, and triggered a cascading series of delays and cancellations for ten days. The root cause wasn't the weather—competitors recovered within days. There was a breach of trust in the data. An operational system was making decisions based on information that no longer reflected reality.

When you look from a distance, these industries—autonomous driving, healthcare, aviation—may seem unrelated. But they are connected by a single principle: When the data fails, the shared world of humans and intelligent systems breaks apart. From that point on, no amount of intelligence—no machine learning model, no predictive algorithm, no clever agent—can save the partnership. Coagency falls apart because one partner sees ghosts.

Understanding Trust in Data

The companies that will dominate the Intelligence Age will not be those deploying the most AI agents. They will be those that create the conditions for *human-machine coagency at scale*. That capability begins with trust in data. Trustworthy data is not a single attribute. It is an architecture built across three tiers: *foundational, operational, and resilient trust*.

TIER 1: FOUNDATIONAL TRUST

Data integrity: accuracy, completeness, and temporal alignment. AI agents act on streams of signals, not static records. When National Grid uses AI to forecast electricity demand, even small timing mismatches across data feeds weaken forecasts and produce recommendations that conflict with operators' control screens. Trust collapses—not because the AI is flawed, but because its

inputs lack temporal coherence. The same logic applies across sectors. At Morgan Stanley, a thirty-second delay in volatile markets means advisors and AI agents see different market realities, leading to conflicting client advice. At Boeing, codesign breaks down if an AI references last week's aerodynamic model while an engineer reviews this morning's wind-tunnel data. Coagency fails when humans and machines operate on different "nows."

Lineage and provenance: the story behind the data. Trust requires traceable reasoning. When Moderna uses AI for vaccine design, no candidate sequence is accepted unless researchers can trace every input—from genomic databases to clinical literature. Provenance allows a scientist to say, "I understand how you reached this conclusion." JPMorgan Chase learned this while rebuilding equity trading systems. Traders resisted AI not because it underperformed, but because decisions were opaque. Trust emerged only after the bank exposed full reasoning chains: which signals mattered, which patterns were invoked, and which risk constraints applied. Coagency is grounded in transparency, not in algorithmic sophistication.

Rights and permissions: legitimacy by design. India's Account Aggregator network embeds trust architecturally. Customers—not institutions—grant, limit, and revoke data access. Permissions travel with the data, allowing AI-driven advisors and underwriting agents to operate on information that is explicitly authorized. GE Vernova faces similar challenges in power generation. Utilities own operational data; GE owns sensor algorithms; service partners need partial access. Fine-grain permissions ensure each party sees what it is entitled to—and nothing more. This is not merely compliance. It is the foundation for cross-organizational coagency.

TIER 2: OPERATIONAL TRUST

Semantic richness: data that describes reality, not just measurements. Treasury Wine Estates initially struggled to use AI for vineyard management because sensor data lacked meaning. "Soil moisture: 23 percent" is insufficient

without context—depth, soil type, recent irrigation, weather conditions, and vine growth stage. Once the data was semantically enriched, coagency emerged. When AI recommended reducing irrigation, managers could see that the system understood the vines’ phenological stage and that mild water stress would improve grape quality.³⁸ Human judgment and machine reasoning converged.

Stability and versioning: predictable evolution. Walmart’s thousands of AI agents rely on shared product hierarchies and customer segments. When one division changes classifications without coordination, downstream agents produce flawed recommendations. To prevent this, Walmart enforces strict versioning of data schemas. Changes are staged, visible, and coordinated. During transitions, both merchandisers and AI agents know which version they are using. Trust comes from predictability, not stasis.

Observability and feedback: continuous calibration. UPS’s On-Road Integrated Optimization and Navigation (ORION) routing system creates trust through bidirectional visibility. When drivers override AI-recommended routes, they provide structured reasons. Both human and machine see the outcome: If the override worked, the model learns; if not, the driver does. Drivers also see why routes were suggested in the first place. This mutual observability transforms AI from a compliance tool into a learning partner.

TIER 3: RESILIENT TRUST

Resilience against drift and manipulation. At Caterpillar, subtle sensor calibration drift degraded maintenance predictions over time. A 2 percent error in hydraulic pressure readings led to misjudged failure timing—small enough to escape detection but large enough to erode technician confidence. Caterpillar responded with continuous calibration verification.³⁹ Prediction errors are monitored, patterns flagged, and recalibration triggered. Technicians are alerted when degradation is detected. Coagency survives because both agents recognize and correct decay together.

Contestability and override rights: partnership, not subordination.

Trust requires the ability to challenge recommendations. At Boeing, engineers retain absolute override authority over AI-suggested structural changes—but overrides initiate dialogue. Engineers document reasoning; models adapt. Both agents learn. Maersk applies similar principles in container shipping. When managers override AI routing, the system probes what the humans saw—weather intelligence, geopolitical risk, or customer constraints not captured in structured data. Overrides are not failures of coagency; they are expressions of it.

Trust calibrated to stakes. Not all decisions require the same level of trust. Spotify users tolerate occasional odd recommendations; trust is ambient and forgiving. In contrast, Merck’s manufacturing AI directly affects drug safety, where errors are nonrecoverable. Enterprises must design differentiated trust architectures. Low-stakes domains can operate with lighter verification and greater autonomy. High-stakes domains require redundancy, human-in-the-loop validation, and exhaustive audit trails. Trust infrastructure must scale with consequences.

The Trust Audit: Questions Leaders Should Ask

The data strategy question leaders face is not “How much data do we have?” but rather “Can our humans and machines see the same world at the same time?” This is the foundational question. If you cannot answer yes with confidence, coagency is impossible. The perception gap will doom every AI initiative, no matter how sophisticated the models.

When they disagree, can both agents trace why? Disagreement is inevitable and often valuable. But if differences cannot be resolved through a shared understanding of data lineage and reasoning, the partnership will collapse into

tribal camps—humans who trust humans, machines that trust machines, and no bridge between them.

If the machine is wrong, can humans safely intervene without breaking the partnership? Contestability is not a failure mode; it's a feature. Organizations that punish overrides or treat them as AI failures will never achieve sustainable coagency. The question is whether your systems are designed to learn from human corrections rather than simply bypassing them.

Do we architect trust differently for ambient versus critical decisions? One-size-fits-all trust requirements either strangle innovation or court catastrophe. The most mature organizations explicitly categorize their AI applications by stakes and build trust architectures appropriate to each level.

Can both agents learn from failures together? This may be the ultimate test. When an AI-human partnership fails—makes a wrong prediction, recommends a poor decision, misses an opportunity—does the failure strengthen the partnership or weaken it? If your answer is “weaken,” your trust architecture is brittle.

Why Trust Fails: The Three Gaps

As we've seen throughout this exploration, trust failures follow recognizable patterns. Understanding these patterns helps leaders diagnose problems before they become catastrophic.

The perception gap: seeing different realities. The phantom braking incident illustrates the first gap. Humans and machines occupied the same physical space but perceived entirely different worlds. The driver saw an empty highway; the vehicle saw an obstacle. When perception diverges, no amount of intelligence can bridge the gap. The partnership dissolves because there is no shared reality upon which to base coordinated action. This gap appears wherever data quality issues create divergent views of the same situation. When National Grid's operators see one demand forecast on their screens while the

AI agent is optimizing for a different forecast—even minutes out of sync—the perception gap destroys coagency.

The explanation gap: understanding the why. The healthcare sepsis prediction system revealed the second gap. Even when the data arrived correctly, clinicians couldn't understand *why* certain patients triggered alerts while others with similar vital signs did not. The machine was making decisions that the humans couldn't trace. Intelligence without explanation is intelligence that cannot be partnered with. JPMorgan Chase closed this gap by building complete lineage visibility into its trading systems. Traders could trace every decision back through the chain of data and reasoning. The explanation gap disappeared, and trust emerged.

The temporal gap: seeing different times. Southwest Airlines' scheduling collapse during the 2022 winter storm demonstrated the third gap. Human schedulers were looking at real-time conditions, while the systems were working from data updated hours earlier in batch cycles. Both were seeing the truth—but they were seeing different moments in time. When humans and machines operate on different clocks, even perfect data becomes untrustworthy because it describes a world that no longer exists.

These three gaps—perception, explanation, and temporal—account for most trust failures in human-machine partnerships. The trust architecture described in this chapter directly addresses each gap: Data integrity closes the perception gap, lineage and provenance close the explanation gap, and temporal synchronization closes the temporal gap.

What We've Learned Working with Clients on Data Trust

We have observed several situations in which data trust is essential. When working with customers, we often start with a seemingly simple question: *Do you trust your data enough to allow a machine to act on it?*

Almost every executive—barring a few—instinctively answers, “Yes.” But as we move deeper into analyzing their systems, processes, and decisions, that confidence begins to soften. The reason is always the same: Trust in data is not based on a feeling. It is an operational reality—one that must be demonstrated, not implicitly assumed.

Over the years, in industries such as healthcare, consumer goods, finance, logistics, and manufacturing, we’ve learned that an organization can’t achieve agentic intelligence until it honestly addresses its trust in data. Trusted data must be precise, complete, consistent, governed, and available precisely when decisions are made—not days, hours, or minutes later. Humans and AI agents work in the world as it *is*, not as it once was. If the data is delayed or conflicted, the partnership falls apart; coagency weakens.

At the same time, we’ve seen how dramatically things change once a company commits to building a solid foundation. The first step is always making the invisible visible—bringing data out of its silos so that every system, team, and agent sees the same version of reality. This unification is rarely glamorous work—we can tell you that; we have scars to prove it! It involves reconciling decades of definitions, resolving contradictions that have been normalized, and compelling the organization to agree on what, for example, a customer, an incident, or an asset is. It’s hard work—but we believe it’s worth the effort. When that shared reality takes shape, the entire enterprise breathes a sigh of relief. Questions stop turning into debates. Decisions become quicker. For the first time, AI systems are beginning to act in ways that seem aligned with human judgment.

The second pattern we’ve observed is that trust isn’t something you achieve once; it’s a moving target. Data must be cleaned, enriched, and refreshed continuously—sometimes multiple times a day. We’ve helped hospitals eliminate duplicate patient records that caused clinical alerts to misfire. We’ve helped insurers reconcile mismatched policy profiles so underwriting models can function clearly. We’ve helped hotels detect silent data drift that was

distorting demand forecasting. In every case, the goal was the same: to enable organizations to rely on their data without questioning whether the information was stale, duplicated, incomplete, or inaccurate.

Governance is the quiet engine behind this reliability. When governance works, people barely notice it. When it doesn't, chaos follows. We've helped companies build lineage systems that show exactly where each data point came from, how it was transformed, who touched it, and which rules it must obey. Executives often underestimate how liberating this transparency can be. Once you can trace your data, you can trust it. Once you can trust it, you can confidently allow AI systems to reason with it.

And finally, we've learned that trust has a tempo. For humans and systems to act together, data must arrive not just clean and consistent but on time. In a modern enterprise, milliseconds matter. A routing agent cannot rely on last night's logistics data. A hospital cannot rely on laboratory results that are updated every four hours. A fraud-detection model cannot wait for overnight batch processing. Real-time intelligence demands real-time truth.

Put simply, we've learned this: *Trusted data is not a technical asset—it's a behavioral one.* It changes how organizations think, decide, and collaborate. It enables AI systems not to replace humans but to stand alongside them, sharing a common reality. And in every client engagement, the same insight emerges: The journey to agentic intelligence does not begin with algorithms, platforms, or models. It starts with trust.

The Data Infrastructure of Intelligence

Leaders often rush to deploy more AI or scale AI across the enterprise. These are the wrong priorities, or at least premature ones.

Instead, ask: *Can your data be trusted enough for AI to act with you, rather than around you?*

If the answer is no, then deploying more AI will only deepen isolation. You will have intelligent agents that cannot partner with humans because they lack a shared, trusted view of reality. You will have automation islands, not agentic intelligence.

If the answer is yes—if you have established the three tiers of trust, strategically prioritized data trustworthiness, and designed distinct trust levels for different decision stakes—then you have set the right tone. You can foster true coagency at scale. You can build an organization where human judgment and machine intelligence genuinely amplify each other.

This is what the Intelligence Age demands. Not sophisticated algorithms alone. Not big data. But trustworthy data that enables humans and machines to see the same world, reason together about it, and act in concert.

Trust quietly and decisively becomes the new infrastructure of intelligence. It is the foundation on which coagency stands. And coagency is the mechanism through which enterprises achieve accurate agentic intelligence—not as isolated pockets of automation, but as coordinated partnerships that multiply capability across the organization.

The companies that understand this are not investing primarily in models or computing power. They are investing in the unglamorous work of data provenance, temporal synchronization, semantic richness, and rights management. They are building the infrastructure of trust. And in doing so, they are building the enterprises that will thrive in an age where intelligence—human and machine, coordinated and aligned—becomes the ultimate competitive advantage.

Trust is the currency. Coagency is the mechanism. Agentic intelligence is the outcome. The transformation begins not with deploying AI, but with earning the trust that makes a human–machine partnership possible.

* * * *

In the Intelligence Age, trust collapses faster than ever. Intelligent systems operate at machine speed, and when they fail—because of flawed data, weak governance, or careless design—the consequences spread instantly. What once unfolded over weeks now happens in seconds: Operational disruption, regulatory scrutiny, and public backlash arrive together. As Warren Buffett famously warned, reputation takes years to build and minutes to lose.⁴⁰ In AI-driven systems, the clock runs even faster. Trust is no longer something companies manage after the fact; it must be engineered into the system from the start.

This reflects a deeper shift. In the Industrial Age, trust was largely reputational—you trusted institutions because of their brand, history, or authority. In the Intelligence Age, trust is behavioral and systemic. People judge organizations by how their digital systems behave: whether data is protected, decisions are explainable, and AI is used responsibly. Consumers increasingly withdraw permission when data is mishandled, and companies that lead in digital trust consistently outperform their peers. Trust has moved from a hygiene factor to a source of advantage—not because it sounds virtuous, but because intelligent systems cannot function without it.

Seen this way, trust is no longer an abstract value. It is the foundation. Doctors will not rely on AI diagnostics if data quality is uncertain. Investors will not follow algorithmic advice if transactions are opaque. Citizens will not accept autonomous systems or digital public services without transparency and accountability. When trust is embedded—through auditable data flows, clear accountability, and governed agency—participation expands, and coagency becomes possible. Agentic intelligence does not earn legitimacy through clever algorithms alone; it earns it by behaving in ways that people can see, verify, and rely on. In the Intelligence Age, trust is not optional—it's the foundation.

And this is the lesson that closes the chapter.

In the Intelligence Age, trust is not what you say; it is what your systems prove. It is not just earned once; it is renewed continuously. It is not built by

marketing; it is built by the data architecture. Coagency—the partnership between humans and machines—depends entirely on this foundation. Without trusted data, machines behave unpredictably, and humans withdraw. With trusted data, humans and machines extend each other's capabilities, confident that they are acting on the same reality.

In the decades ahead, the companies that thrive will not be the ones with the most data, the largest models, or the flashiest AI. They will be the ones that treat trust as a design principle, embed it in their organizational workflows, and make it visible in their actions. Trust is no longer a virtue. It is a competitive advantage. It is the connective tissue that allows intelligence—human and artificial—to flourish.

And it begins, unavoidably, with data that is trustworthy.

OceanofPDF.com

SECTION 2

INVEST TO MONETIZE

Given a 10 percent chance of a 100 times payoff, you should take that bet every time.

—JEFF BEZOS

In the first pillar, Strategic Intent, we referenced the parable of the blind men and the elephant—a story that illustrates how organizations can mistake parts of data for the whole truth. Strategic intent is about seeing the full picture, not just isolated pieces. But seeing, as leaders quickly learn, is only the beginning. Insight without action is like a mirage—beautiful from afar but unable to satisfy your need.

To understand this second pillar—Invest to Monetize—imagine a caravan crossing a desert. The travelers know their destination and the landscape they want to reach. But vision alone can't move them forward. As temperatures rise and the terrain becomes more challenging, the group needs to refill its water supply. Scouts return with news: Three wells lie ahead.

The first well is nearby, shallow, and dependable—there's enough water to support the caravan for another day. The second is farther away, deeper, and less certain. Reaching it requires taking a detour, exercising patience, and having the courage to invest time and resources without an immediate reward. The third is even farther, beyond the visible ridge, mentioned more in faith than proof. If it exists, it could change not just the caravan's progress but also its possibilities.

A caravan that drinks only from the nearest well survives the present but never reaches its imagined future. A caravan that learns to draw from all three—today's certainty, tomorrow's potential, and the future's promise—is the one that turns a journey into a destination.

This is the core principle of the first rule in this pillar: Organizations must learn to invest by allocating scarce resources across different time horizons. It is a discipline of maintaining the present while building the foundation for what's next. Many companies struggle here because Horizon 2 (H2)—the middle ground where old processes are phased out and new ones are built—is the most challenging terrain. It requires not only capital but also conviction.

But the caravan's story does not end in the desert. When the travelers reach the final hill, expecting more sand, they instead see a port city emerging along the coast, illuminated by a lighthouse whose beam cuts across the sea. With this single structure, the entire rhythm of commerce has shifted. Ships no longer wait for daylight. Goods move steadily through fog, wind, and uncertainty. Navigation becomes safer, faster, and more predictable. The lighthouse is not the commerce itself—but it is the intelligence that enables new commerce.

This embodies the core of the second rule: Intelligence gains value not through hindsight but through forward-looking action. When organizations invest in systems that reveal what is likely to happen—rather than just what has happened—they unlock new opportunities to create and capture value.

Even the brightest lighthouse cannot guide a ship if sailors do not know how to interpret its beam. Therefore, the caravan eventually discovers the final truth embedded in this pillar: Investing in human capital is integral to the journey. Agentic intelligence does not diminish the importance of people; it amplifies it. It encourages them to collaborate with machine agents, integrate probabilistic signals into their decisions, and navigate confidently in a world where recommendations and actions are increasingly cocreated. This is the third rule: Prepare people for coagency.

Together, these three rules establish the logic of Invest to Monetize. They demonstrate that intelligence isn't something you buy; it's something you develop gradually through resource allocation, architectural discipline, organizational learning, and human development. Strategic Intent offered us a new perspective. This pillar provides a new way to build. The wells, the lighthouse, the sailors, and the caravan all convey one message: To monetize intelligence, you must invest in it—over time, through systems, and among people.

In the upcoming chapters, we examine each rule—how companies balance investment across different time horizons, how they turn intelligence into new types of value, and how they train their people to succeed in an era where humans and machines reason together. Think of this as your invitation to go beyond vision and engage in the practical, organizational, and human efforts required to make intelligence profitable for your business.



Strategic Intent

Invest to Monetize

Implement to Win



RULE 5

INVEST ACROSS HORIZONS TO BUILD AGENTIC INTELLIGENCE

The riskiest thing we can do is just maintain the status quo.

—BOB IGER, DISNEY

In every business cycle, there comes a moment when leaders must face an uncomfortable truth: The strategy that made their organizations successful is no longer the one that will carry them forward. Today, that moment is here, marked by the rise of agentic intelligence—the coordinated partnership between humans and machines that enables an organization to sense, decide, and act with exceptional coherence.

What surprises us—as a CEO and an academic—is not the power of AI but how leaders are responding to it. Across industries, pilots flourish; they show promise. AI models outperform traditional analytics. Copilots assist knowledge workers. Predictive maintenance tools prevent breakdowns. Digital twins simulate complex systems. There is no shortage of technological promise in pilots and experiments.

And yet, ask any CEO what frustrates them most about AI implementation, and the answer is always the same: *We can build pilots. We cannot scale them.*

This gap—between what is technically possible and what is organizationally achievable—is the defining leadership challenge during this phase of transition to the Intelligence Age. It’s fundamentally an *investment problem*. Not in money alone but in the allocation of four kinds of capital: financial capital, human capital, digital capital, and attention capital.

Most companies continue to invest as they did during the industrial era: heavily in the present, minimally in the future, and almost not at all in transformative change. In such an environment, agentic intelligence cannot take root. It requires more than automated workflows or modern platforms. It requires organizational adaptation, in which human and digital agents learn to work together.

The problem is not that companies lack ambition or vision. They lack what we call the “multihorizon discipline.”

The Horizon 1-2-3 Discipline

For decades, the industrial era trained leaders to think in one-year cycles, evaluate projects individually, and optimize for efficiency. In an era where intelligence flows through data, that logic is woefully insufficient. Agentic intelligence requires unified data foundations, continuous learning loops, and cross-functional interoperability—all of which demand multihorizon investment.

Horizon 1 (H1) is the present, the world of execution—automated workflows, operational dashboards, and improved forecasting. These investments deliver ROI next quarter and build credibility. But they create a subtle trap: They reinforce the logic of the current business rather than building the architecture of the next one.

Consider a retailer implementing AI-powered demand forecasting. The project delivers clear value, but the forecasting system improves efficiency without questioning whether the entire fulfillment model needs to be reimaged. The organization believes it's "doing AI" when it's just automating the status quo.

Horizon 3 (H3) is the world of imagination: autonomous vehicles, intelligent factories, agent-driven commerce, and agentic healthcare that's personalized and preventative. These investments inspire teams, generate patents, attract top talent, and create powerful demos. Organizations love H3 because they're aspirational. But in most organizations, H3 innovations remain disconnected experiments—proof that the future exists but not the specific proof that the company can achieve it.

The autonomous robot performs well in the lab, but its capabilities are limited if it can't connect to the warehouse management system. The diagnostic agent achieves impressive accuracy but lacks access to real patient records, which are distributed and protected under different privacy protocols. The digital twin produces fascinating simulations but has no pathway to influence actual operations.

***Horizon 2 (H2)** is where the failure to scale becomes visible.*

Why H2 Is the Crucible of Agentic Intelligence

H2 is where workflows need to change. This is where legacy systems must be phased out, where data needs to be unified across domains and regions, where data must flow from internal sources to ecosystem partners, and where governance shifts from control to coordination. This is where roles, responsibilities, and decision rights must evolve to support intelligent agents working alongside humans.

Consider a global manufacturer deploying predictive maintenance across fifty factories in fifteen countries. The H1 pilot at one German factory is working perfectly. The H3 vision of fully autonomous manufacturing is inspiring. But H2 is where reality intrudes: The Polish factory uses different sensor protocols. The Chinese facility can't export data to the cloud for regulatory reasons. The Mexican plant relies on paper-based maintenance logs. Each facility defines *machine downtime* differently and has its own escalation process.

To scale, the company needs to standardize data schemas, build secure cross-border pipelines, digitize paper processes, integrate with multiple equipment application programming interfaces (APIs), establish new decision rights (Who can shut down a machine based on AI recommendations?), create new skills (maintenance technicians interpreting AI insights), and update metrics (How do you measure a breakdown that was prevented?).

None of this is glamorous. Most of it is organizationally painful. All of it is necessary—absolutely essential—for agentic intelligence.

This is also where humans must learn to work *with* intelligent systems, not just operate them. The maintenance technician needs to understand when to trust the AI's recommendation and when to override it. The plant manager needs to coordinate between human judgment and algorithmic forecasts. The supply chain planner needs to cocreate solutions with AI agents that have access to data that humans cannot process. This is not automation—it is augmentation. It is coagency.

H2 is where governance shifts from control to coordination. When an AI system recommends a course of action, who is responsible? When human judgment conflicts with algorithmic recommendations, who decides? When an AI agent initiates a transaction, which budget does it draw from? These aren't technical questions—they're organizational design questions requiring sustained leadership attention.

You've seen the headlines: Most companies have active AI pilots, but few have achieved scaled impact. The reason is clear to us—and we're sure to you as well: Models are available, but organizational systems aren't built to incorporate them. In every survey, the same obstacles appear—fragmented data, lack of cross-functional ownership, insufficient governance, and underinvestment in the architecture needed to redesign how humans and machines collaborate.

This is why H2 is the crucible of agentic intelligence. Unfortunately, it's also why H2 is the most underfunded horizon.

The Four Capitals That H2 Demands

It demands *financial capital*—because workflows, data platforms, and legacy systems must be rebuilt rather than patched. But H2 financial investment differs from H1 investment. H1 projects have clear, measurable ROI. H2 investments deliver benefits that compound over time and span multiple use cases.

When a bank unifies customer data across divisions, the immediate ROI is elusive. But that unified data becomes the foundation for hundreds of future AI applications. The value is not in the data platform itself, but in what it enables: humans and AI agents working from a shared understanding of reality.

It requires *human capital*—new skills in agent orchestration, prompt engineering, and knowledge modeling. But these aren't just technical skills. Organizations need people who can translate between business operations and AI capabilities, who understand both algorithmic logic and the psychology of change management, who can help their colleagues learn to work *with* AI rather than compete against it.

The hardest human capital challenge is cultural: helping people shift from seeing AI as a threat to seeing it as a collaborator, from defending old ways of working to cocreating new ones. As Doug McMillon emphasized at Walmart,

the goal isn't just to get AI right but to “create the opportunity for everybody to make it to the other side.”⁴¹

It demands *digital capital*—because systems must shift from static repositories to dynamic intelligence networks where humans and machines share context. In the Industrial Age, digital capital meant stability. In the Intelligence Age, it means systems that learn through use, adapt to new data sources continuously, and enable both human experts and AI agents to access the same operational truth.

It also requests *attention capital*—because none of this can succeed without sustained executive focus. H2 investments don't clamor for attention the way operational crises do. They move slowly, require patience, and demand understanding of technical complexity. But without sustained executive attention, H2 initiatives stall, and the organization continues accumulating AI pilots that never enable true human–AI collaboration.

H2 also requires leadership to perform a dual mental task: advancing the present while envisioning the future. Leaders must simultaneously protect today's margins while funding tomorrow's capabilities, optimize current operations while designing new workflows, and deliver quarterly results while investing in multiyear platforms. This cognitive load is immense. It requires the courage to invest billions in infrastructure that may not show clear ROI for years.

Learning from H2 Leaders

What distinguishes leaders who successfully fund H2? Three consistent patterns:

First, they reframe investment questions. Instead of “What's the ROI?” they ask, “What capability are we building for humans and machines to collaborate?” They understand that some investments create optionality—the

ability to move quickly as the H3 future becomes clearer and weak signals strengthen.

Second, they protect H2 budgets from quarterly pressures. They create separate funding pools or shield H2 investments from normal ROI scrutiny applied to H1 projects. This recognizes that different horizons require different evaluation frameworks.

Third, they translate technical complexity into a strategic narrative. They help boards and employees understand why data unification matters for human–AI collaboration, why workflow redesign is strategic, and why governance frameworks enable coagency between human experts and intelligent agents.

Let's examine how specific leaders have executed this pattern.

MERCEDES-BENZ: REIMAGINING HUMAN–MACHINE PARTNERSHIP

Mercedes-Benz CEO Ola Källenius has been explicit: “The car is turning into a digital device. We must master software, or we will not lead the next generation of luxury.”

Källenius is clear-eyed about the challenge: “The electric drivetrain has a higher variable cost than the internal combustion engine. We have more than a technical challenge: we have an economic challenge. We need vertical integration in software. That's something we need to own.”

H2—the shift to software-defined vehicles, unified vehicle data, and AI-powered autonomy—will determine competitiveness. “We are the architects of our own operating system, a chip-to-cloud architecture,” Källenius explains. The Mercedes-Benz Operating System (MB.OS) platform represents this H2 bet: unglamorous, expensive, organizationally disruptive, but essential.⁴²

The vision transforms the human–vehicle partnership. Mercedes-Benz's fleet becomes “computers-on-wheels connected to the cloud,” with real-time data continually gathered within the car and then fed to the cloud to enhance

safety, convenience, and personalization. Källenius disclosed that the company generated €1 billion in software-enabled revenue in 2022 and expects “high single-digit billion euros” by decade’s end. “With MB.OS,” he says, “a Mercedes will actually improve with age.”

The driver and the vehicle become collaborative agents—the human provides judgment and preferences, and the machine provides real-time data processing and anticipatory assistance.

MICROSOFT: BUILDING THE ENTERPRISE INTELLIGENCE LAYER

Before OpenAI became widely known in late 2022, Satya Nadella was already designing Microsoft’s H2: moving Microsoft Office to the cloud, unifying enterprise productivity data, acquiring LinkedIn, developing Azure as an intelligence hub, and creating the Microsoft Graph.

“More than any transformation before it, this generation of AI is radically changing every layer of the tech stack, and we are changing with it,” Nadella wrote in his 2025 shareholder letter.⁴³ Azure exceeded \$75 billion in revenue with 34 percent growth, and Microsoft Cloud surpassed \$33 billion in quarterly revenue.

When generative AI matured, Microsoft was uniquely positioned, not because it had the best models, but because it had built the only enterprise-scale intelligence layer capable of hosting them. “We’ve moved from talking about AI to applying AI at scale,” Nadella emphasized at Microsoft Build 2024. “By infusing AI across every layer of our tech stack, we’re winning new customers.”⁴⁴

The H2 architecture enables millions of knowledge workers to collaborate with Copilot across Microsoft 365, GitHub, and Teams. Over one hundred million monthly active users work alongside AI agents, which act not as replacements but as collaborators. This is agentic intelligence at enterprise scale.

WALMART: 2.1 MILLION HUMANS LEARNING TO WORK WITH AI

Doug McMillon at Walmart offers perhaps the most striking example of multihorizon leadership. “It’s very clear that AI is going to change literally every job,” McMillon warned in late 2024. “Maybe there’s a job in the world that AI won’t change, but I haven’t thought of it.”⁴⁵

What distinguishes McMillon’s approach is refusing to separate technical transformation from human investment. “We love what technology can do,” he said at CES 2024, “but we’re building it in a way that creates better careers at the same time. It creates better customer experiences and a stronger business.”

Walmart expects to hold headcount flat at 2.1 million over three years while revenue grows through AI-enabled productivity. “Our goal is to create the opportunity for everybody to make it to the other side,” McMillon said. The company partnered with OpenAI to launch certification programs, created a new “agent builder” role, and rolled out its My Assistant AI tool to fifty thousand associates.

But the real H2 work is in how humans and AI collaborate in the supply chain. Suresh Kumar, Walmart’s global CTO, explained: “A modern supply chain requires built-in intelligence that can forecast what customers want and when, and orchestrate movement of different products.” Walmart built an industry-leading forecasting system using patent-pending machine learning, incorporating dozens of data types—from historical sales to weather forecasts to social media trends.

The result? Humans and AI agents working together: the humans providing business judgment and customer insight, the AI processing signals that no human could integrate. Walmart wasn’t chasing a model. It was funding a future in which human associates and machine agents collaborate seamlessly to serve millions of customers daily.

JPMORGAN CHASE: 250,000 EMPLOYEES LEARNING TO COCREATE WITH AI

Jamie Dimon at JPMorgan Chase represents the most ambitious experiment in human–AI collaboration at scale. “AI is real,” Dimon said in 2024. “There will be no job, no process, no function that won’t be affected by AI—mostly for the positive.” But Dimon is clear-eyed about what this means: “We took AI and data out of technology. It’s too important. We put AI at the management table. Dr. Manuela Veloso, head of JP Morgan Chase AI research, reports to our president and me. Are we doing enough? Are we doing it right?”⁴⁶

To Dimon, this isn’t a technology project. It’s an unprecedented organizational transformation. JPMorgan spends \$2 billion annually on AI. He believes that the immediate benefits that they have seen so far are only the tip of the iceberg. He’s bullish on what the future might hold for the bank and the financial industry.

The numbers tell the story of human–AI collaboration at scale: 250,000 employees now have access to the bank’s LLM Suite platform. Roughly 150,000 use it weekly to summarize reports, conduct research, and analyze contracts. The platform generates investment banking decks in thirty seconds—work that previously required teams of analysts working through the night.

But here’s what makes this an H2 investment rather than simple automation: Humans aren’t being replaced; they’re being augmented. The analyst who once spent hours creating presentations now uses that time for strategic thinking and client relationships. The risk officer who once manually reviewed contracts now focuses on complex judgment calls. The customer service representative who once struggled with edge cases now has an AI copilot suggesting solutions.

Dimon is realistic about the challenge: “People shouldn’t put their head in the sand. It is going to affect jobs. There’ll be fewer jobs in certain functions.” But the company is “aggressively retraining and redeploying” talent, investing in human capital alongside digital capital.

As Dimon told a data conference, “The head of AI is at every single meeting I have with management teams so that the bank can stay on top of things. We have investments in 100 different companies. We’re testing and learning.”⁴⁷

This is what H2 looks like at the world’s largest bank: not humans *versus* machines but humans *and* machines learning to collaborate, cocreate, and coevolve. JPMorgan is building what its chief analytics officer calls a “fully AI-connected enterprise”—not to eliminate human judgment but to augment it at every level.

Dimon’s vision? “Your children are going to live to 100 and not have cancer because of technology. And literally they’ll probably be working three-and-a-half days a week.” That optimistic vision requires significant H2 investment today in the unglamorous tasks of data unification, workflow redesign, governance frameworks, and helping hundreds of thousands of people learn to work effectively with intelligent agents.

MAERSK AND SHELL: INDUSTRIAL INTELLIGENCE AT GLOBAL SCALE

Maersk recognized that global supply chains were entering an era in which data velocity and cross-border intelligence matter more than fleet size. It modernized its entire data infrastructure, invested in real-time visibility platforms, integrated IoT sensors across containers, built digital twins of ports, and developed predictive analytics for route and fuel-use optimization.

Every investment was an H2 bet: unglamorous, expensive, organizationally disruptive. But the payoff was transformational: Maersk became capable of orchestrating multiagent systems where humans and AI collaborate to automate scheduling, routing, and customs processing. Its intelligence advantage came from years of sustained investment in systems allowing humans and machines to work together.

Shell took a similar path. Instead of waiting for the perfect AI model, Shell invested in a digital ecosystem supporting continuous learning and automated decision-making. It funded digital twins of refineries and platforms, retooled workflows for predictive maintenance, built subsurface AI capabilities for reservoir modeling, and centralized operational data to support agents making real-time safety recommendations alongside human operators.

As Shell's CTO expressed: "We're building the digital foundations now so that autonomy becomes possible later."⁴⁸ But that autonomy isn't about removing humans—it's about enabling humans and machines to collaborate on problems too complex for either to solve alone.

Architecting Coagency Across Horizons

Across Walmart, JPMorgan, Maersk, Shell, Mercedes-Benz, and Microsoft, a familiar storyline emerges: They didn't try to "do AI." They invested in becoming intelligent enterprises where humans and machines work together.

Doing AI is a technical activity. Becoming intelligent is a structural metamorphosis that enables coagency. These companies understood the core tenets:

- Algorithms are cheap. Architecture for human–AI collaboration is not.
- Pilots are easy. Workflows enabling cocreation are hard.
- Insights are abundant. Trust between humans and AI is scarce.
- Talent can be hired. A culture of collaboration must be built.
- Data can be collected. Shared understanding must be engineered.
- Models can act. Organizations must adapt to enable coagency.

But what exactly makes their H2 investments different from traditional technology investments? Three distinctions matter.

H2 INVESTMENTS BUILD SHARED CONTEXT, NOT JUST INFRASTRUCTURE

Traditional IT investments develop systems operated by humans. H2 investments build environments where humans and AI collaboratively share operational context.

Consider the difference at JPMorgan. A traditional investment banking system stores deal data and provides reports. JPMorgan's LLM Suite doesn't just store and report—it develops shared understanding with the human analyst. When an analyst asks it to generate a pitch deck, it understands the client's context, the competitive landscape, the regulatory environment, and the deal structure. It doesn't replace the analyst's judgment; it provides enough context to offer useful suggestions that the analyst can refine.

This requires fundamentally different architecture. The system needs access to the same information the human has—and the human needs to understand what information the AI has access to. This mutual transparency enables trust. Without it, the human either over-relies on AI recommendations (dangerous) or ignores them entirely (wasteful).

At Mercedes, this shared context means the vehicle knows not just where the driver is going, but how they prefer to get there, what driving style they're comfortable with, and what assistance they typically need. The car and driver develop a shared understanding over time. This isn't automation—it's partnership.

H2 INVESTMENTS ENABLE CONTINUOUS LEARNING LOOPS, NOT ONE-TIME SOLUTIONS

Traditional projects have defined scopes and endpoints. H2 investments create systems that improve through use—and where both humans and machines learn together.

At Walmart, the supply chain forecasting system doesn't just predict demand. It learns from every decision a category manager makes. When the manager overrides a recommendation, the system learns why. When the manager accepts a recommendation that proves correct, the system gains confidence. When the manager accepts a recommendation that proves incorrect, both the system and the manager learn together.

This is fundamentally different from traditional decision support systems, which provide recommendations but don't learn from human responses. The learning loop requires infrastructure: data capture of human decisions, feedback mechanisms, retraining pipelines, and—most importantly—time for humans to understand how the system learns so they can teach it more effectively.

Shell's predictive maintenance systems work similarly. When a human operator overrides a shutdown recommendation because they know something the AI doesn't (perhaps recent maintenance not yet in the system), the AI learns to incorporate that signal. When an operator follows a recommendation that prevents a failure, both parties learn what indicators matter most. This continuous loop means the system gets better *because* humans work with it, not despite human involvement.

H2 INVESTMENTS REQUIRE NEW GOVERNANCE

Traditional governance centralizes decision rights. H2 governance coordinates decisions across human and machine agents.

At JPMorgan, when an AI agent flags a potentially fraudulent transaction, who decides whether to block it? The AI has pattern recognition; the human expert has contextual judgment. The customer is time-sensitive; the regulatory framework has compliance requirements. H2 governance establishes protocols for coordinating these different forms of intelligence.

This isn't about "humans in the loop" versus "humans out of the loop." It's about architecting appropriate collaboration for different decisions. Some

decisions benefit from AI acting autonomously with human oversight. Others benefit from AI providing recommendations that humans execute. Still others benefit from real-time collaboration where humans and AI iterate together.

Traditional governance models don't accommodate this complexity. They assume hierarchical decision rights: Either the system decides, or the human decides. Agentic intelligence requires what we might call "distributed decision rights"—frameworks that specify when different agents (human or machine) have primary authority, when they must consult, when they can act independently, and how conflicts get resolved.

This is what Shell's CTO meant by "building the digital foundations now so that autonomy becomes possible later." The autonomy isn't complete machine control—it's the ability for humans and machines to make collaborative decisions with appropriate independence for each party, depending on the situation.

The Rule in Practice

Let's be specific about what leaders do in H2, because the gap between aspiration and execution is where most organizations fail.

They unify data semantics across silos. Not just centralizing data but establishing shared definitions. When Walmart's AI forecasts demand, and the category manager reviews it, and the supply chain planner acts on it, all three—two humans and one AI—need to understand demand the same way. This requires months of unglamorous work reconciling definitions across systems, regions, and business units.

They redesign workflows to enable collaboration. When JPMorgan gives 250,000 employees access to AI assistants, they don't just deploy software—they aim to redesign how work flows. The analyst no longer starts by building slides from scratch; they start by specifying objectives to the AI. The AI doesn't just generate output; it explains its reasoning so the analyst can refine it. This

workflow redesign requires studying how work actually happens, identifying collaboration points, and creating new processes.

They invest in explainability and transparency. When Mercedes-Benz's vehicle suggests a route, the driver needs to understand why. When Shell's system recommends shutting down a refinery process, the operator needs to understand the reasoning. This requires investing in interfaces, visualizations, and explanation mechanisms that traditional AI systems don't provide. It's not enough for the AI to be right—it must be interpretable by its human partner.

They create new roles and career paths. Walmart's agent builder role didn't exist five years ago. These people help merchants design and deploy AI agents tailored to their specific needs. They're not data scientists, and they're not merchants—they're translators who understand both domains and can help humans and machines work together effectively. Creating such roles requires reimagining career development, compensation structures, and organizational reporting lines.

They establish feedback and learning mechanisms. When JPMorgan's AI suggests a recommendation the analyst rejects, that rejection becomes data. When Walmart's AI forecast proves to be wrong, that error becomes training data. When Mercedes-Benz's driving assistant makes a suggestion that the driver ignores, that interaction becomes a learning signal. But capturing this feedback requires instrumentation, storage, processing, and—most importantly—ensuring humans know their interactions help the system improve.

They build trust gradually. None of these companies implemented agentic intelligence across the entire organization simultaneously. They began with specific use cases, teams, and workflows. They demonstrated value. They learned what worked. They refined their approach. They addressed concerns. They celebrated successes. They acknowledged failures. Only after this iterative process did they scale.

This is why H2 takes years and millions of dollars, depending on the enterprise's size. It's not deploying models—it's re-architecting how

organizations work when intelligence becomes distributed across humans and machines.

The leaders of these companies—McMillon, Dimon, Källenius, Nadella—demonstrated a common trait: an ability to think with two minds—and two time horizons—at once.

One mind is fully anchored in the present: protecting margins, serving customers, mitigating risk. Another mind is relentlessly building the architecture of collaboration between humans and machines. This duality requires discipline, political courage, and tolerance for uncertainty. It requires a willingness to fund projects that executives cannot yet fully articulate. It requires betting on the long arc of intelligence rather than immediate optimization.

Leaders who succeed at this adopt a different internal narrative. They don't ask: *What is the ROI of a future we cannot yet see?* They ask: *What investments must we make today so humans and machines can work together tomorrow?*

Agentic intelligence will reward only those leaders who understand that intelligence is not a feature or function. It's a capacity—a capacity for humans and machines to sense, decide, and act together—that must be built through sustained, multihorizon investment.

Agentic intelligence is the destination. Investment across horizons is the vehicle.

And H2—the hard middle where humans learn to collaborate with machines—is the road almost no one wants to travel, but everyone must.



RULE 6

MONETIZE BY EMPOWERING AGENTS TO ACT

Knowing is not enough; we must apply. Being willing is not enough; we must do.

—LEONARDO DA VINCI

In early 2024, a Rolls-Royce Trent XWB engine operating on a long-haul flight registered a subtle deviation in pressure readings deep within the turbine. The signal was minor—well within certified operating limits—but it followed a pattern that the engine’s health-monitoring system had learned to treat as an early warning.

That signal was automatically sent to Rolls-Royce’s operations and analytics system, where it was compared to historical and fleetwide performance data from millions of engine flight hours. The system detected an emerging problem, adjusted the expected maintenance window, and allowed maintenance plans to be modified ahead of time—so parts, expertise, and resources could be prepared to match the aircraft’s schedule. The flight

continued smoothly. There were no delays, no disruptions, and no visible issues for the airline or its passengers.

What mattered was not that data was collected or that an insight was generated. What mattered was that the system acted.

This is Power by the Hour in the Intelligence Age. Under this longstanding service model, airlines do not pay Rolls-Royce for engines as hardware; they pay for availability and uptime. The economic value is realized not when a report is generated but when an intelligent system—operating within human-defined boundaries—detects risk early, coordinates maintenance, and prevents failures before they become operationally visible. In this model, the engine itself acts as an agent: sensing conditions, updating predictions, and triggering coordinated responses across engineering, maintenance, and supply networks.

At the core of this capability is what Rolls-Royce describes as a continuous, bidirectional data thread. Data flows from the engine to analytics platforms and back again, combining real-time sensor data with maintenance records, operational context, and environmental conditions. These systems constantly update life-limit projections for critical components and enable engineers to concentrate monitoring on specific parameters when necessary. This approach goes beyond traditional predictive maintenance, which often detects issues too late or operates at coarse intervals. Here, degradation is tracked at the individual-engine level, allowing interventions well before scheduled maintenance cycles.

The lesson isn't just about aviation. Monetization isn't derived solely from prediction, data, or insight alone. It results from coordinated action—systems that can sense, decide, and respond quickly enough to change outcomes. Power by the Hour succeeds because intelligence is embedded directly into operations, not just layered on top as analysis.

This chapter explores a principle that many leaders are only beginning to grasp: *In the Intelligence Age, value is created and captured not by knowing more, but by acting sooner—and acting together.*

Four Monetization Patterns

For decades, organizations believed value came from analysis. Better dashboards. Deeper analytics. More accurate predictions. The entire business intelligence (BI) industry was built on this premise: Give leaders better insight, and they will make better decisions.

But the Intelligence Age has redefined the economics of value creation. In a world operating at machine speed, insight alone is no longer enough. Monetization flows not from knowing but from *doing*—from intelligent systems capable of taking contextual, responsible, and timely action on behalf of people and organizations.

Consider the differences:

- *Insight* indicates that an engine part will likely fail after five hundred flight hours. *Action* preorders the replacement, schedules the maintenance window, coordinates the spare engine, and updates the fleet optimization model—all before the airline's operations team finishes their morning coffee.
- *Insight* tells a logistics company that a shipment route is suboptimal. *Action* reroutes the shipment, updates delivery estimates, notifies the customer, and adjusts pricing for future similar routes—while the package is still in motion.
- *Insight* tells a bank that a customer's portfolio is misaligned with their goals. *Action* rebalances the portfolio within preapproved parameters, documents the rationale, and sends a summary the customer can review—before market conditions change.

The value gap between these scenarios is enormous. Insight creates the potential for value; action realizes it. And in markets where milliseconds matter, where opportunities expire before humans can respond, where complexity exceeds human processing capacity, the organizations that can act

intelligently in real time will monetize, while those that merely analyze will watch.

We've identified four monetization patterns, each representing a different way that action creates economic value.

PATTERN 1: FROM SELLING ASSETS TO SELLING UPTIME

Rolls-Royce pioneered this model in 1962—yes, over sixty years ago—with the original Power by the Hour, but the Intelligence Age has transformed what uptime means. It's no longer just about availability—it's about continuous optimization through intelligent action.

Today's Rolls-Royce engines don't just fly; they adapt. They optimize fuel efficiency based on weather patterns. They adjust performance parameters based on altitude and load. They coordinate with ground systems to minimize turnaround time. The TotalCare program uses real-time engine data analytics to optimize maintenance scheduling, reducing unscheduled repairs and extending service intervals by approximately 25 percent.⁴⁹

An Airbus A350 generates 2.5 terabytes of data daily from six thousand sensors.⁵⁰ However, data generation alone does not create value—action based on intelligence does. When the engine autonomously adjusts variable stator vane actuators using performance data to prevent failures before they happen, that action produces monetizable value. Similarly, when the system predicts component replacement needs and initiates supply chain actions, that also creates value. Proactive maintenance, which reduces aircraft on-ground time from days to hours, further generates value.

The business model has shifted from transactional (selling an engine) to ongoing (selling intelligent uptime). Airlines pay for the *actions* that intelligent systems take to keep aircraft flying, not just for the systems themselves.

PATTERN 2: FROM SELLING SERVICES TO SELLING OUTCOMES

Ocado Group has transformed grocery fulfillment into an agentic orchestration problem. Its warehouses aren't facilities—they're ecosystems where thousands of robots collaborate to fulfill orders at impossible speed.

At Ocado's customer fulfillment centers, an AI air traffic control system orchestrates fleets of robots moving at four meters per second across 3D storage grids. The system communicates with each robot ten times per second, coordinating movements with five-millimeter precision. In 2024, Ocado's On Grid Robotic Pick system picked over thirty million items, completing fifty-item orders in five minutes—a task that took over an hour using conventional methods.⁵¹

But here's what makes this agentic: The robots aren't following programmed routines. They're making continuous decisions. AI algorithms determine optimal routes for thousands of simultaneous picks. Machine vision systems identify grasp points on unfamiliar products. Reinforcement learning enables robots to adapt when items are positioned unexpectedly. Behavior cloning allows the fleet to learn from human operators, then share that knowledge globally.

As Ocado explains: "Our robots use Machine Learning to pick and pack products with no prior understanding of what they are and how they should be handled. The unique benefit is that the arm can adapt on its own and share these learnings with robots across the globe, so each error becomes a lesson for the entire fleet."⁵²

Ocado doesn't sell robotic equipment—it sells continuous fulfillment intelligence as a service. Retailers pay for outcomes: fulfilled orders, optimized inventory, reduced waste. The monetization comes from intelligent actions happening thousands of times per minute across the facility, not from the hardware itself.

And critically, this is human–machine collaboration at scale. Warehouse staff aren’t replaced—they’re augmented. Robots handle repetitive picking while humans manage exceptions, oversee quality, and coordinate operations. The system extends picking hours beyond human working hours, enabling redeployment of staff to higher-value activities. Action creates value; humans and machines together optimize that value creation.

PATTERN 3: FROM SELLING DATA TO SELLING INTELLIGENCE-IN-ACTION

Thomson Reuters represents this pattern in professional services. Its integration of legal intelligent agents—trained on Harvey AI—doesn’t just give lawyers access to information. It enables them to act faster and more comprehensively than humanly possible.

The transformation isn’t subtle. Previously, a junior associate might spend forty hours researching case law for a brief. With Harvey integrated into workflows, that same research happens in minutes. But the value isn’t in the speed of data retrieval—it’s in what the lawyer can now *do*: Take on more clients, deliver richer analysis, anticipate counterarguments, and expand their practice’s scope.

Harvey doesn’t replace the lawyer; it amplifies their capacity to act. The attorney still applies judgment, crafts strategy, and advises clients. But they do so with an intelligent assistant that can sift through millions of legal documents, draft preliminary briefs, identify relevant precedents, and suggest analytical frameworks.

This pattern appears across professional services. McKinsey’s Lilli, trained on decades of consulting knowledge, doesn’t sell insight—it enables consultants to act on insights faster. Medical diagnostic agents don’t just identify conditions—they recommend treatment protocols and coordinate care teams. Financial planning agents don’t just analyze portfolios—they execute rebalancing within client-approved parameters.

The monetization shifts from selling access to knowledge (subscriptions, hourly rates) to enabling intelligent action (outcomes delivered, cases won, diagnoses confirmed, portfolios optimized). Law firms using Harvey effectively charge for better outcomes achieved faster. Consulting firms with AI augmentation deliver insights that translate immediately into client action.

PATTERN 4: FROM SELLING INSIGHTS TO EMBEDDING INTELLIGENCE IN CUSTOMERS' LIVES

Mercado Libre demonstrates how agentic intelligence creates entirely new markets by enabling action at scale. Across Latin America, the platform's machine agents approve microloans in seconds, orchestrate fraud detection across millions of transactions, and optimize fulfillment flows in real time.

These aren't back-office improvements—they're revenue-generating actions that enable commerce for millions of small businesses that previously couldn't participate in the digital economy. A street vendor in São Paulo can accept digital payments, access working capital, and manage inventory because intelligent agents handle the complexity. The vendor acts (sells products), the agents act (process payments, assess credit risk, prevent fraud), and Mercado Libre monetizes the coordinated action.

Similarly, in industrial systems, Siemens and Schneider Electric are turning energy intelligence into subscription businesses. Their digital twins and optimization agents don't just report on building performance—they automatically balance consumption, detect faults, and adjust HVAC, lighting, and power systems. Building managers aren't buying devices; they're buying continuous intelligent action that reduces energy costs and improves occupant comfort.

In healthcare, agentic care loops are emerging, in which intelligent systems remind patients to take medications, adjust care plans based on real-time sensor data, and notify clinicians when intervention is needed. The monetization aligns with value-based care: reduced readmissions, improved

outcomes, lower total cost of care. The agents don't replace clinicians—they extend clinical capacity by handling routine monitoring and coordination, freeing clinicians to focus on complex decision-making.

Across these patterns, a common economic logic emerges. *Agentic intelligence monetizes by collapsing the gap between knowing and doing.*

First, action compounds. Insight is a point-in-time artifact. Action creates momentum. When Rolls-Royce's engine adjusts performance parameters thousands of times per flight, each micro-optimization compounds into fuel savings, extended component life, and increased reliability. One insight about optimal cruise altitude is valuable. Ten thousand real-time adjustments based on that insight create measurable economic value.

Second, action scales beyond human capacity. A skilled technician can monitor dozens of machines. An intelligent agent can monitor thousands. A portfolio manager can actively manage fifty client accounts. An intelligent agent can rebalance thousands of portfolios simultaneously based on each client's unique parameters. The monetization comes from doing what humans cannot do at scale, not from knowing what should be done.

Third, action enables new business models. Power by the Hour, a new revenue model, was made possible by intelligent engines that enable continuous optimization. Ocado's fulfillment-as-a-service model requires robotic orchestration that no human team could coordinate. These aren't improvements to existing business models—they're entirely new ways of creating and capturing value that only become possible when intelligence can act.

Think Coagency, Not Automation

But here's what distinguishes true agentic monetization from simple automation: These systems succeed because humans and machines act together.

At Rolls-Royce, engineers define maintenance protocols, but agents execute them at fleet scale. At Ocado, warehouse managers oversee operations, but robots handle thousands of simultaneous picks. At Thomson Reuters, lawyers craft legal strategy, but Harvey handles research and drafting. At Mercado Libre, risk analysts set fraud parameters, but agents process millions of transactions.

This is not humans versus machines. It's humans and machines creating value together through coordinated action—what we call coagency. The human provides judgment, context, ethics, and strategic direction. The machine provides scale, speed, pattern recognition, and tireless execution. Together, they achieve outcomes neither could accomplish alone. The monetization comes from human-machine collaboration, not from replacing humans with machines.

This distinction matters economically. Automation reduces costs. Coagency creates new revenue. Automation optimizes existing processes. Coagency enables new business models. Automation is a one-time efficiency gain. Coagency creates continuous value through learning loops in which humans and machines improve together.

If monetization begins with action, the next question is uncomfortable: How must organizations change so that agents—human and digital—can act together coherently?

Most organizations can analyze. Their teams can strategize. Their dashboards can visualize. But when it comes time to *take action*, they falter—not because they lack intelligence, but because they're designed for a previous era's constraints. Workflows, approvals, incentives, and governance were designed for a world in which only humans acted, at human speed, with human limitations.

Agentic intelligence introduces a different cadence: continuous sensing, continuous reasoning, continuous action. This rhythm is incompatible with

workflows depending on weekly meetings, quarterly reviews, or hierarchical decision chains.

Thomson Reuters learned this with Harvey. In law firms where Harvey sits outside the workflow—a tool lawyers can optionally consult—adoption is slow, and value creation remains marginal. In firms where Harvey is integrated into workflows, where its outputs become first drafts that lawyers refine, value emerges quickly. Lawyers take on more cases, meet deadlines earlier, and deliver richer analysis.

The technology didn't change. The workflow redesign allowed revenue to surface.

AIA Group in Asia illustrates the same pattern in insurance. Wellness agents can nudge customers toward healthier behaviors, but these agents create value only when integrated into policyholder journeys, claims handling, and risk modeling. When workflows remain siloed, nudges become noise. When workflows are unified, agents reduce claims, improve retention, and enhance lifetime value.

Humans must relinquish certain microdecisions to digital machines. Agents must escalate complex judgments to humans. And workflows must choreograph their interaction. When designed well, this choreography becomes a revenue engine. When neglected, no amount of intelligence yields value.

As workflows evolve, business models must evolve with them. Agentic intelligence doesn't just automate processes—it reshapes the economics of what companies sell. In traditional business models, monetization was tied to discrete transactions: a machine bought, a flight booked, a loan approved. But agentic workflows create continuous, ongoing value.

This pattern reveals a larger truth: Agentic intelligence monetizes not by selling predictions but by embedding intelligence into the flow of the customer's life or business. Companies that persist in selling insights—or

worse, selling access to dashboards—will be left behind. Insight does not compound. Action does.

The organizations that thrive will be those that (a) redesign workflows such that humans and agents share labor fluidly, (b) build trust architectures that allow agents to act responsibly, (c) align business models with continuous value loops, and (d) create economic structures where every action—however small—drives value forward.

Three Questions Leaders Must Ask (and Answer)

Monetizing agentic intelligence is not merely a technical challenge. It is a leadership challenge that requires redefining work, reshaping incentives, retiring legacy workflows, and giving both humans and machines the authority to act. Leaders must answer three hard questions:

First: What actions create the most value for our customers? Not what insights do they need, but what actions improve their outcomes. Rolls-Royce didn't ask, "What data do airlines want?" They asked, "What actions keep aircraft flying?" Ocado didn't ask, "What information do retailers need?" They asked, "What actions fulfill orders faster?"

This question forces brutal honesty about where value lives. A bank might find that customers value monthly portfolio reviews (insight) less than automatic portfolio rebalancing when market conditions create opportunities (action). A manufacturer might realize customers don't want predictive maintenance reports—they want their equipment to never fail in the first place.

Second: Where can intelligent agents act autonomously within the boundaries we set? This involves identifying decisions that are (a) time-sensitive (benefit from immediate action), such as market trades, inventory replenishment, and route optimization; (b) rule-based (can be governed by

clear parameters), such as credit approvals within limits, automated refunds, and standard maintenance procedures; (c) pattern-based (benefit from processing more data than humans can), such as fraud detection, quality control, and resource allocation; (d) repeatable (occur frequently enough to justify automation), such as customer service inquiries, scheduling, and order fulfillment; and (e) monitorable (can be audited and overridden when necessary), meaning all agent actions must be transparent and reversible.

The key is defining guardrails that enable action while preserving human oversight. A treasury agent might execute foreign exchange hedges up to \$10 million autonomously but escalate larger transactions. A legal research agent might draft preliminary briefs but flag complex precedential questions for attorney review. An energy optimization agent might adjust building temperatures within a two-degree range but require approval for larger changes.

These guardrails aren't constraints on intelligence—they're enablers of trust. By clearly defining when agents can act and when they must consult, organizations create the confidence necessary for widespread adoption.

Third: How do we redesign workflows such that humans supervise actions rather than execute them? This is the most difficult question because it requires changing how people work, what they are held accountable for, and how they are measured. It means shifting from “Did you do the work?” to “Did the work get done well?” It involves measuring outcomes instead of activities.

In a manufacturing plant, computer vision agents now inspect every part coming off the line. The quality inspector no longer spends their time sampling, measuring, and recording results. Instead, they supervise the system: reviewing edge cases the agents flag, identifying systemic issues, and refining inspection parameters so defects are caught earlier and more consistently. Their role is not eliminated—it is elevated. Quality improves across the entire line, not just in sampled batches.

In a law firm, AI tools such as Harvey handle initial research and draft preparation. Junior associates no longer spend most of their time searching case law or assembling first drafts. Their work shifts to validating reasoning, contextualizing precedents, and shaping legal strategy. They handle more cases, learn faster, and develop judgment earlier—while the firm serves clients more effectively.

In a hospital, intelligent agents automate routine coordination—scheduling follow-ups, sending reminders, and tracking compliance. Clinical coordinators focus instead on complex cases: intervening when patterns indicate risk, coordinating among specialists, and tailoring care plans for patients who require the most attention. Outcomes improve, and human expertise is applied where it is most effective.

Across all three settings, the pattern is the same. Humans move from executing tasks to supervising intelligent action, from managing transactions to shaping outcomes, from doing the work to ensuring the work is done well. That is the practical meaning of agentic intelligence.

But this transition requires more than new technology. It requires:

New metrics. Organizations must measure what matters in an agentic world: outcomes delivered, problems prevented, opportunities captured, improvements achieved. Traditional productivity metrics (tasks completed, hours worked, cases handled) become misleading when agents handle volume and humans handle quality.

New incentives. Compensation systems must reward people for optimizing agent performance, not just their personal output. The quality inspector should be rewarded for improving agent accuracy. The associate should be rewarded for extracting maximum value from Harvey. The coordinator should be rewarded for the health outcomes their supervision enables.

New skills. People need to learn agent supervision: understanding how agents make decisions, recognizing when agents need human intervention,

improving agent performance through better training and feedback. This is distinct from traditional technical skills or domain expertise—it's a new capability that combines technical understanding with domain knowledge and judgment.

New culture. Organizations must create environments where human-machine collaboration is celebrated, where asking agents for help is seen as smart rather than weak, where improving agent performance is recognized as a valuable contribution. This requires leadership explicitly modeling coagency.

Two Journeys into Agentic Monetization

For all the talk about AI and machine intelligence, the hardest truth is this: Many companies are still standing at the starting line, weighed down by fragmented data, siloed processes, and workflows that were built for a slower age. Some companies face this reality not metaphorically but literally—because all their customer data lives in the hands of one person.

That was the situation at Schneider Electric, a global leader in energy management. With more than 128,000 employees and operations across 115 countries, Schneider powers homes, hospitals, data centers, and factories—yet for years, it hadn't paid adequate attention to data modernization.

As the company grew, leadership began to see a strategic irony: Schneider had become world-class at orchestrating electrons but had almost no ability to orchestrate *its own customer signals*. There were no unified identities, no consistent records, and no reliable intelligence about who bought what, who needed what, or how customers navigated the company's sprawling offerings. The firm knew it needed not just more data but the ability for data to act—to power workflows the way electricity powers a grid.

So, Schneider invested in a new cloud-based data platform to become the single source of truth for five million organizations and thirteen million individuals. More than twenty systems—Salesforce CRM, SAP ERP, Oracle

Fusion Cloud Financials, quoting systems, e-commerce platforms—were consolidated into a unified, continuously refreshed intelligence layer. What had once taken days or weeks became nearly instantaneous.

“In less than five minutes, a new customer entry is certified,” recalls Loïc Tordo, Schneider Electric’s director of master data. “That helps our staff be more efficient and provides a better service experience to our customers.”⁵³

What Schneider discovered next is the essence of agentic monetization: Once high-quality data begins to flow, *value begins to accumulate everywhere*.

The company uncovered millions of “lost” customers—accounts incorrectly marked inactive because siloed systems couldn’t recognize them. Suddenly, entire pools of demand reappeared, expanding Schneider’s pipelines without a single new advertising dollar. Logistics costs fell as address data became more reliable, misshipments declined, and reshipments slowed. Sales teams could finally see the full scope of a customer relationship, rather than operating blindly with partial histories.

Schneider wasn’t yet deploying fully autonomous agents, but it had established the most important prerequisite: a data foundation that could someday enable agents to act on behalf of the company and its customers. It was the first step toward monetizing intelligence.

The second story shows what happens when that foundation translates into revenue.

Republic Services presents a revealing case of how an operations-first company began reframing itself around customer intelligence. As the second-largest environmental services company in the US, Republic had long excelled at what mattered most in its industry: trucks, routes, safety, and efficiency. Customer intelligence, by contrast, was secondary. As one internal leader put it, Republic was “a Fortune 300 company no one recognizes unless they pick up your trash.”

That mindset was embedded in its systems. Heather Tweedy, director of technology, described it bluntly: “If there’s an invoice, that’s an account. And

we don't need to think bigger than that.”⁵⁴ Billing defined identity. Relationships were fragmented. A single customer overseeing dozens of sites appeared in systems as dozens of disconnected records. The company could run operations at scale—but it could not truly see its customers.

The inflection point came after Jon Vander Ark became CEO in 2021 and articulated a long-term ambition to elevate customer experience as a strategic priority. At the same time, Republic was expanding through acquisitions that brought stronger customer analytics capabilities into the portfolio. The challenge became unavoidable: Either unify customer data across the enterprise or constrain future growth.

Republic initially attempted to build its own customer data foundation internally. The effort proved fragile and difficult to scale. The company ultimately pivoted to a modern data platform designed to unify customer identity across dozens of operational and transactional systems.

When the new platform went live, the scale of the problem—and the opportunity—became visible. As Tweedy recalled, “Our billing system had no sense of person. If one person oversaw 15 accounts across 30 sites, the system treated them as multiple different identities. When those went into the new platform, they became one clean person. Sales could finally see the relationship. It was huge.”

This was more than data cleanup. It changed how the business could act. Once Republic could see customers holistically—what they bought, where they operated, and how services connected across locations—cross-sell and upsell opportunities surfaced naturally. Revenue opportunities that had been invisible in fragmented systems became actionable. Republic publicly discussed ambitious growth goals tied to improved customer intelligence and exceeded early expectations, prompting leadership to raise those targets and deepen its investment in data.

The lesson is not that Republic suddenly became a digital native. It is that agentic revenue depends on seeing customers as connected entities, not

disconnected transactions. When data is unified around real relationships, intelligent systems can support sales, service, and growth in ways that operational excellence alone never could.

Republic's experience underscores a broader truth of the Intelligence Age: *You cannot act intelligently on customers you cannot see—and you cannot see them without connected, trusted data.*

Republic's journey demonstrates what Schneider's story hinted at: *You cannot monetize agentic intelligence until you can see. But once you can see, you can act—and acting is what generates revenue.*

Both companies began with data foundations, but their long-term trajectories point toward something more transformative. Schneider now has the infrastructure to deploy autonomous energy-management agents that optimize customer consumption. Republic is positioned to build routing agents, pricing agents, safety agents, and customer-engagement agents. They are stepping into the early stages of the agentic world.

In that world, data is not merely a record of the past. It's the fuel for the human and AI agents shaping the future.



RULE 7

INVEST IN AGENTIC LITERACY TO BUILD SHARED CONTEXT

The illiterate of the twenty-first century will not be those who cannot read and write, but those who cannot learn, unlearn, and relearn.

—ALVIN TOFFLER

In December 2021, as the world was still navigating the pandemic, Moderna made an unusual announcement. The biotech company behind one of the first COVID-19 vaccines revealed it was launching an AI Academy—not for a select group of data scientists, but for every employee. CEO Stéphane Bancel’s message was direct: “We believe AI is a key enabler of our ability to build the best version of Moderna now and into the future. This AI Academy will enable us to make AI part of the company’s ways of working, part of our DNA.”⁵⁵

The AI Academy wasn’t designed to turn biologists into programmers. Scientists would learn how AI could advance drug design. Manufacturing teams would understand how algorithms could optimize production. Regulatory staff would grasp how models could accelerate approvals. But the deeper ambition was about something subtler: building shared context—a

common language and understanding between humans and intelligent systems that would allow genuine collaboration rather than wary coexistence.

Four years later, that bet seems foresightful. We're moving into an era where AI agents don't just analyze—they act. They coordinate logistics, draft contracts, modify pricing, approve requests, and oversee portfolios. The obstacle to creating value is no longer the complexity of models but whether humans and machines can perceive the same world clearly enough to collaborate effectively.

This is where most organizations are stuck today. Impressive pilot projects stall at scale. Agents perform well in the laboratory but falter in production. Employees work around the systems rather than with them. The systems are smart—but they're smart in their own world.

Rule 7 addresses this gap. It argues that escaping “pilot purgatory” requires a fundamental shift in how organizations think about workforce capability. For the past decade, enterprises have invested heavily in data literacy—teaching people to read dashboards, interpret key performance indicators (KPIs), and spot trends in reports. That movement was necessary. It helped companies stop flying blind. But it was designed for a world where humans make decisions and systems merely inform them. That world is disappearing.

In the agentic world now emerging, agents don't just inform decisions—they make them. They route trucks autonomously, adjust supply chains dynamically, rebalance portfolios in real time, approve loan applications, draft legal contracts, and schedule workforce deployments. The workflow has fundamentally changed. Agents continuously sense and interpret their environment and propose or execute actions within guardrails, while humans supervise, override, refine, and reframe. The question is no longer “Can you read the dashboard?” It's “Can you collaborate with a system that sees the situation differently than you do—and still reach coherent decisions together?”

That requires a different capability entirely: *agentic literacy*. And building it requires investment across four types of organizational capital—financial,

human, digital, and attention—to create environments where humans and agents share sufficient context to work as genuine partners. The payoff isn't theoretical. It's the difference between pilots that impress and transformations that monetize.

Here is the mistake most organizations make: They try to jump from dashboards to autonomous agents without doing the hard work in between. You cannot move from H1 to H3 without crossing H2—the unglamorous, capital-intensive work of building shared context and learning how to supervise intelligent action.

H1 gave organizations data literacy: the ability to read metrics, interpret dashboards, and make better human decisions. H3 promises autonomous agents that act at scale. But H2 is where value is actually unlocked. It is where shared context is operationalized, where systems learn to make their reasoning visible, where humans learn to work with agents as teammates rather than tools, and where both sides develop a common understanding of goals, constraints, and edge cases. When H2 is underfunded, H3 never scales.

Most companies know how this ends. They maintain H1 out of necessity and invest aggressively in H3 out of fear—while starving H2. The result is familiar: impressive demos, stalled pilots, and expensive AI that sits unused or quietly bypassed.

The human capability required in H2 is *agentic literacy*, and it cannot be taught in a classroom. It is learned through practice.

First, it requires *delegation and orchestration*, not just interpretation. Data-literate managers spot anomalies and escalate them. Agentic-literate managers set objectives, supervise how agents pursue them, refine constraints, and intervene only when context is misread. This is learned the same way people learn to manage others—through coached, hands-on experience.

Second, it demands *continuous adaptation*, not static training. Agentic systems evolve constantly. Learning must be embedded in daily work, with

feedback flowing in both directions as humans refine agents and agents reshape human judgment.

Finally, it requires *coagency*, not tool use. Traditional software forced humans to adapt. Agentic systems require mutual adaptation. Humans contribute judgment, ethics, and contextual nuance; agents contribute speed, scale, and pattern recognition. Trust is calibrated over time through repeated interaction.

Without H2—and without agentic literacy at scale—organizations don't fail dramatically. They stall quietly, stuck between dashboards they've outgrown and agents they can't yet trust.

Why Pilots Fail to Scale

One pattern is consistent across industries. A logistics company deploys an AI-powered routing system that reduces delivery times by 15 percent in a controlled pilot. However, when it attempts to expand it across regions, dispatchers quietly override the system's suggestions. When the system recommends skipping a trusted local partner in favor of what appears to be a more efficient route, workers hesitate. Not because they believe the AI is wrong—they simply don't understand what it's seeing or what it values. One dispatcher sums up the issue well: "It's not that the system is dumb. It's smart. But it's smart in its own world."

Industry surveys consistently show poor ability to scale AI pilots to enterprise-level deployment. The issue is rarely the underlying model's performance. It's the human infrastructure surrounding the technology. The primary bottleneck is a lack of agentic literacy at scale—the workforce isn't prepared to integrate, govern, and trust autonomous agents in mission-critical operations.

This is an H2 failure. Organizations bought the technology but didn't invest in the capacity to absorb it. They tried to jump from H1 to H3 without

building the bridge. Escaping pilot purgatory requires systematic investment in four interconnected areas.

First, siloed data and governance structures must give way to architectures that provide agents with trusted, accessible, governed data streams across the enterprise. Agents require enterprise-wide visibility; robust governance builds the trust necessary for autonomy. Without this, agents operate with partial information and humans can't verify their reasoning.

Second, a lack of trust and adoption can't be solved with formal training courses. Organizations must shift to embedded, on-the-job coaching and pair programming with AI. Literacy grows only through daily practice, feedback loops, and seeing value in real workflows. When Anthropic studied how its own engineers work with AI, it found that roles were fundamentally shifting—from writing code to managing AI agents. One engineer estimated their work had shifted “more than 70 percent to being a code reviewer/reviser rather than a net-new code writer.” This transformation doesn't happen in a classroom.

Third, undefined roles and risk allocation keep organizations trapped in low-risk pilots. Clear human-agent teaming playbooks must specify who—human or agent—is accountable for what decisions, when escalation is required, and how responsibility shifts as agents gain capability. Without this clarity, pilots never graduate to production because no one is willing to take accountability for autonomous actions.

Fourth, static metrics designed for individual productivity become misleading in a world of human-agent collaboration. New KPIs must measure human-agent team productivity—metrics such as agent intervention rate, goal completion time, and the quality of human feedback that improves agent performance. The focus must shift from individual output to the performance of the augmented team.

This isn't a nice-to-have training initiative. It's a hard-headed strategy that determines whether technology investments deliver measurable ROI. It ensures that when you deploy agents that learn in the loop (Rule 8), act on shared

representations (Rule 9), and are governed in motion (Rule 10), they operate in an environment where humans can guide, correct, and amplify them.

When Investment Becomes DNA

Let's return to Moderna for a moment. The AI Academy wasn't a symbolic gesture. It was a multiyear, multicapital commitment to build agentic literacy from the ground up. The partnership with Carnegie Mellon wasn't just about credentialing—it was about understanding the shared context. Moderna's leadership understood that if only a few experts knew how the models worked, everyone else would treat them either as opaque oracles or bureaucratic hurdles.

So they made a comprehensive bet. Financial capital came in the form of a substantial, ongoing commitment to partner with an elite university and build comprehensive curricula. Human capital meant asking busy professionals—scientists in the lab, engineers on the floor—to commit real time to learning, not just attend a lunch-and-learn. The digital capital involved integrating AI tools into everyday research and operations so learning connected to actual work. And attention capital meant putting this initiative near the center of the company's transformation narrative, with CEO ownership, not buried in an HR footnote.

The curriculum covered data quality, visualization, statistical thinking, machine learning algorithms, and crucially, AI ethics. Not to turn biologists into data scientists but to help them understand what AI could do for their problems and what its limitations were. The result isn't a company where AI runs everything. It's a company where scientists, engineers, and business leaders can reason with their AI agents. They grasp, at a high level, how models see relationships in data. They know where to be skeptical and where to lean in. And they can explain to the systems what constraints truly matter—safety,

manufacturability, regulatory implications—beyond what’s encoded in a loss function.

Moderna used data literacy as a starting point. But its real bet was on agentic literacy: building a workforce that can collaborate with AI, not just use it.

Scale and Depth

If Moderna represents depth, Walmart represents scale. For a company employing over two million people worldwide, the temptation might be to keep AI expertise centralized—a team in Bentonville building tools for everyone else. Instead, in 2025, Walmart announced a partnership with OpenAI not just to create AI-first shopping experiences but to embed AI training across its entire workforce.

The customer-facing element made headlines: Shoppers would soon shop inside ChatGPT, with AI planning meals, restocking essentials, and handling instant checkout. But the more significant commitment was providing free, customized AI training and certification to all US associates starting in 2026—part of nearly \$1 billion invested in skills training through 2026.

The training will be delivered through Walmart Academy, the company’s massive internal training platform serving 3.5 million participants, tailored to the day-to-day realities of Walmart’s operations. How frontline employees can work with AI-powered scheduling. How supply chain staff can collaborate with inventory agents. How associates can use conversational AI to solve customer problems.

Walmart framed it clearly: “We know the future of retail won’t be defined by technology alone—it will be defined by people who know how to use it. By bringing AI training directly to our associates, we’re putting the most powerful technology of our time in their hands—giving them the skills to rewrite the playbook and shape the future of retail.”⁵⁶

This isn't fully agentic literacy yet—some of it still falls under data literacy and tool usage. But it's laying the critical groundwork for H2. When associates soon work with AI-powered scheduling agents, inventory copilots, and customer-facing tools, they'll need more than comfort with technology. They'll need to understand what the agents see, where they lack context, when to intervene, and how to correct them. Walmart is building that capacity at scale, with the same comprehensive investment pattern: financial capital in the platform and training, human capital in asking millions of associates to build new skills, digital capital in weaving AI into product catalogs and operations, and attention capital in CEO-level messaging that positions AI as amplifying jobs rather than replacing them.

When Banks Become Tech Companies

For a more mature example of H2 investment paying off at scale, look at the Development Bank of Singapore (DBS). DBS began its digital transformation in 2014, long before *agentic AI* entered the vocabulary. But the bank understood early that transformation wasn't just about technology—it was about people.

CEO Piyush Gupta led a bold cultural transformation. The bank used the acronym GANDALF to highlight its goals: Google, Amazon, Netflix, Apple, LinkedIn, Facebook—with DBS aiming to join that elite group. But ambition alone doesn't develop skills. DBS made significant, ongoing investments in its people. It became the first Singaporean bank to introduce a professional conversion program to re-skill employees in eight roles. It created Digify, a modular learning pathway that trains employees in agile practices, big data, and journey thinking. It launched DBS Academy, which conducts over ten thousand training sessions each year. It trained eight thousand employees in a program called Data Heroes, teaching how data can be used and what questions it can answer. During COVID-19, it ran a program called Deep

Racer, teaching employees how to program autonomous vehicles—not because DBS runs cars, but to build understanding of how agents learn and behave. The bank formed a Data Chapter, bringing together seven hundred data professionals to scale AI and machine learning value. By 2024, over five thousand employees had received bespoke training across seven hundred courses tailored to data professionals.

The results are tangible. DBS has deployed over 350 AI use cases across consumer banking, institutional banking, and support functions. It accelerated the time to value for AI initiatives from twelve to fifteen months to two to three weeks. Over two years, AI's economic impact on DBS's business more than doubled, from S\$150 million to over S\$370 million, with projections of S\$1 billion in AI-driven revenue within five years.

But the deeper transformation is cultural. DBS created an environment where employees across the organization—not just technologists—can speak the language of data and AI. They understand what models can do, where they struggle, how to frame problems for agents, when to trust outputs, and when to intervene. Crucially, DBS didn't just train people on AI—it changed how people work. It created a new operating model called Managing Through Journeys, where cross-functional squads work together on end-to-end customer experiences, with data professionals embedded within each squad to accelerate AI deployment.

This is H2 investment at enterprise scale: building the human and digital infrastructure that enables shared context. The bank isn't just deploying agents—it's creating an organization where humans and agents genuinely collaborate.

Knowledge Work at the Frontier

If Moderna represents technical depth, Walmart scale, and DBS transformation maturity, then McKinsey and Boston Consulting Group (BCG) represent something else: knowledge work at the frontier. These are

firms whose entire business model is intellectual labor—the very work that AI is poised to transform most dramatically. Both have made bold, enterprise-wide commitments to agentic literacy that go far beyond tentative pilots.

McKinsey deployed Lilli, a proprietary AI platform, across the entire firm. The research reveals a striking gap: Employees are using AI three times more than their leaders expect. While C-suite executives guessed that only 4 percent of employees use AI for substantial portions of their work, the actual figure was 13 percent. This perception gap matters. It reveals that workforce transformation is often already underway—employees are adapting faster than organizations realize—but without the infrastructure, governance, and shared vocabulary to do it well.⁵⁷

McKinsey responded by investing heavily in what it calls “AI literacy, AI adoption, and AI domain transformation,” building baseline fluency, integrating tools into workflows, and developing domain-specific use cases that enhance competitive advantage. Importantly, McKinsey emphasizes that AI adoption is fundamentally a change management challenge, not just a training rollout. Its research on large-scale transformations shows lasting adoption occurs when employees understand what to do differently, believe in why it matters, feel supported by leadership, and see reinforcement in their systems. This is H2 infrastructure—not just teaching people to use tools but redesigning how work gets done.

BCG took an even more aggressive approach. BCG’s philosophy is “early and often”—it understood from day one that generative AI would be transformational, and that to help clients transform, it needed to transform itself first. BCG rolled out ChatGPT Enterprise to all thirty-two thousand employees globally. Not to a pilot group. Not to select practices. To everyone. It also built proprietary tools, including Deckster for generating presentation slides, and GENE, a conversational AI that evolved from a podcast cohost into a client engagement tool.

The results are striking. BCG employees have created over thirty thousand custom generative pretrained transformers—personalized AI applications built through natural language to handle everything from improving prompts to scheduling meetings across time zones to generating creative client talk tracks. Deckster has been used to create or edit slides more than 450,000 times since its March 2024 launch.

This approach has generated what BCG executives call increased “employee joy”—people aren’t just more productive, they’re more satisfied with how they work. The ability to offload administrative tasks, experiment with new approaches, and operate at a higher level of abstraction has fundamentally shifted how consultants experience their work.⁵⁸

The key insight here is that BCG views agentic literacy not as a defensive safeguard but as a competitive advantage. If your business advises clients on transformation, you need to be further along the maturity curve than they are. You can’t credibly guide clients through H2 if you’re still stuck in H1 yourself. Both McKinsey and BCG reveal a vital pattern: Organizations at the forefront aren’t asking whether they should invest in H2 infrastructure—they’re competing on how effectively they build and deploy it.

Agentic literacy isn’t optional. It requires deliberate investment in shared context, common practices, continuous learning, and an honest reckoning with both opportunities and tensions.

The Currency of Coagency

All of this—the training programs at Moderna and Walmart, the transformation journeys at DBS, the enterprise-wide AI adoption at McKinsey and BCG—points to something more fundamental: the central role of shared context in making agentic intelligence actually work.

Imagine a scenario. It’s 2028. A global logistics company has deployed a sophisticated dispatch agent that routes thousands of shipments daily,

optimizing for on-time delivery, fuel costs, and capacity utilization. The agent makes dozens of decisions per minute, nudging the network into balance. On paper, it works beautifully. Delivery times are down. Fuel costs are lower. Exceptions are caught earlier.

But supervisors are quietly overriding the system's suggestions. Dispatchers are working around it. When the agent recommends skipping a trusted local partner in favor of a seemingly more efficient route, people balk. Not because they think the AI is wrong—they just don't understand what it's seeing or what it cares about.

In the Intelligence Age, the main bottleneck is rarely model quality or compute power. It's the gap between how machines see the world and how humans do. Agentic intelligence—the coordinated action of humans and machines together—only emerges when those worlds overlap enough to build shared context.

Shared context means three things. First, humans understand what the agents see. Not in esoteric technical detail but enough to hold a useful mental model. “This routing agent prioritizes on-time delivery and fuel cost, but it doesn't know about our informal agreement with that regional partner—so I need to watch that route.” When people talk about agents this way—as colleagues whose strengths and blind spots they understand—agentic literacy is taking root.

Second, agents understand what humans see. Agents need access not just to data, but to context: the tacit knowledge, soft constraints, relational dynamics, edge cases, and strategic priorities that humans carry but that don't appear in transaction logs. This requires deliberate investment—bringing frontline teams into agent design, capturing their decision heuristics, encoding their exceptions, creating feedback channels where humans can say “technically correct but contextually wrong,” and having that update future behavior.

Third, both parties share a common representation of the problem space. Through data architectures, knowledge graphs, semantic layers,

documentation, and cross-functional dialogue, both agents and humans can see the same entities, relationships, events, constraints, and goals. When this shared representation exists, agents can explain their reasoning in terms humans recognize, humans can guide agents using concepts the systems understand, and both can learn from each other's mistakes.

Without shared context, agents cannot safely take on meaningful tasks, humans cannot trust agents enough to let them act, governance becomes reactive—explaining failures instead of preventing them—and monetization stalls after pilots because scaling requires people who know how to work with agents. With it, something different becomes possible.

Think about Netflix. The company has long called itself “data-informed.” Its analytics teams help make decisions about content choices, product features, recommendations, and artwork. Netflix engineers create data products—telemetry, metrics, dashboards, models—designed with specific users and feedback loops. This focus on data as a product moves toward shared understanding. Content executives, product managers, and data scientists work together, looking at the same metrics, experiments, and patterns and discussing what they mean for decisions: renewing a series, launching a new feature, or changing the sign-up process.

In an agentic Netflix of the future, where agents might propose marketing campaigns, adjust pricing dynamically, or personalize experiences across modalities, this shared context becomes even more crucial. A model won't be trusted merely because it performed well in training. It will be trusted because people understand how it sees the world—and can correct it when that view diverges from reality.

The Four Capitals

The investment pattern across these examples clarifies what building agentic literacy requires. Financial capital is the most visible—companies write checks

to universities, technology partners, and platform vendors. They sponsor academies, boot camps, and new roles. There's a temptation to think the job is done once the budget is approved. But without corresponding investments in human capital, the money yields far less than it could.

Human capital means more than hiring data scientists. It means upskilling the existing workforce to play new roles: supervising agents, refining objectives, spotting context gaps, integrating machine insights into workflows. It means giving people time, not just tools, to learn.

Digital capital is the foundation that makes context shareable. It includes integrated data architectures, knowledge graphs, semantic layers, and observability systems that allow both humans and agents to see the same entities, relationships, and events. A firm with siloed systems can train local models, but it will struggle to deploy coherent agents across operations.

Finally, and perhaps most importantly, there is attention capital. What do senior leaders talk about, ask about, reward, and model? If AI is treated as a side project or a risk to be contained, it will never attract the shared engagement agentic literacy requires. If it's publicly framed as central to the firm's future—and if leaders themselves demonstrate curiosity and openness in learning how agents work—the organization gets a very different message.

These aren't independent investments. They reinforce each other. Financial capital without attention capital results in costly shelfware. Human capital without digital capital leaves people trained on capabilities they can't access. Digital capital without human capital builds sophisticated infrastructure that no one knows how to use. Attention capital without the others is empty signaling.

The Leadership Mandate

Let's be direct about what this requires from top management.

First, shift the focus from data literacy to agentic literacy. Data literacy taught employees to interpret historical information. Agentic literacy teaches them to command the organization's future—to work with systems that act autonomously, to supervise agents as teammates, to foster dynamic coagency. This isn't a rebrand. It's a different capability requiring different investments.

Second, fund H2 infrastructure aggressively. This means building data architectures that provide shared context, not just storage; creating training programs embedded in daily workflows, not episodic courses; establishing new metrics that measure human-agent team performance, not just individual throughput; and defining clear playbooks for human-agent teaming, including responsibility, accountability, and escalation. The temptation is to fund H1 (current operations) and H3 (moonshot projects) while starving H2. Resist it. H2 is the bridge that determines whether your investments in autonomous agents deliver value.

Third, treat agentic literacy as a strategic differentiator rather than a learning and development initiative. The core challenge in scaling AI from pilots to enterprise-wide operations isn't the technology's capability—it's humans' ability to supervise and orchestrate autonomous systems. Organizations that master human-AI collaboration will build formidable moats. Those that don't will watch their expensive pilots gather dust while competitors that invested in literacy pull ahead.

Fourth, model curiosity about how agents work. If senior leaders treat AI as a black box or delegate it entirely to technical teams, that signals to the organization that it's someone else's problem. If leaders ask questions, share examples of working with agents, discuss failures openly, and demonstrate a willingness to learn—even when it's uncomfortable—the organization follows.

Finally, make agentic literacy visible in talent decisions. Hire for it. Promote for it. Reward managers who build team capability in human-agent collaboration. Celebrate employees who find better ways to work with agents,

who surface context gaps, who improve how humans and machines understand each other.

The skills that created your current organization won't sustain it in the Intelligence Age. The winners will be organizations that treat building agentic literacy as seriously as they once treated building operational excellence or customer focus.

Why This Matters for Monetization

This rule sits within the Invest to Monetize pillar because no amount of modeling sophistication can substitute for shared context in coagency. Without it, the other rules in this book can't deliver value. You can't monetize action (Rule 6) if humans don't trust agents enough to let them act. You can't learn in the loop (Rule 8) if feedback channels between humans and agents don't exist. You can't govern in motion (Rule 10) if people can't interpret what agents are doing in real time. You can't build shared representations (Rule 9) if your organization lacks the literacy to understand them.

The early years of AI deployment reveal this pattern clearly. Many organizations have impressive prototypes in labs. Far fewer have systems deeply embedded in workflows, delivering measurable economic value. The difference is often not the algorithm. It's the readiness of the humans around it.

The competitive race isn't just about who has the smartest models. It's about who builds organizations where humans and machines can genuinely work together. That capability doesn't emerge from algorithms alone. It requires deliberate, sustained investment in the literacy that makes coagency possible—financial capital, human capital, digital capital, and relentless attention capital deployed together, one capability, one team, one workflow at a time.

The investments that brought you to H1 won't carry you to H3. Without the bridge of H2—without agentic literacy at scale—your autonomous agents

will remain trapped in pilot purgatory, your transformation budget will yield minimal ROI, and your competitors who invested in readiness will leave you behind.

The companies making these investments now—building the literacy, creating the shared context, developing the capability for genuine human-machine collaboration—are the ones that will monetize agentic intelligence at scale. The rest will continue to run impressive demos while their competitors run the business.

[OceanofPDF.com](https://oceanofpdf.com)

SECTION 3

IMPLEMENT TO WIN

In theory, there is no difference between theory and practice. In practice, there is.

—UNKNOWN

Many organizations today stand at a curious threshold. They have sharpened their strategies, defined their architectures, and made courageous investments in the future. Their pilots are promising, their prototypes sparkle, and their CEOs speak fluently about AI on earnings calls. From a distance, it would seem they are ready to leap into the Intelligence Age with confidence.

Yet something happens when they move from ambition to reality. A system trained on millions of historical records hesitates in front of a live anomaly. A customer service agent receives an AI suggestion that feels tone-deaf to the human nuance of the moment. A pricing model performs beautifully in simulation, then stumbles when competitors shift unexpectedly. Human supervisors override automated actions not because the system is wrong, but because they can't see how it reached its conclusion. A brilliant generative model produces insights no one feels confident implementing.

The organization has intelligence—but not yet *agency*. It can analyze but cannot act fluidly. It can recommend but cannot adapt coherently. It can automate but not coordinate.

This is why the third pillar exists. This is where the AI edge shows up in financial reports.

If the first pillar (Strategic Intent) taught us to reimagine what intelligence should look like in a company, and the second pillar (Invest to Monetize) showed how to assemble the economic and organizational foundations for that intelligence, then this third pillar (Implement to Win) reveals what it takes to make intelligence *work under real conditions*. It deals not with aspiration but with friction; not with architecture but with behavior; not with the promise of AI but with the lived experience of deploying it.

To understand this pillar, picture not a machine or a factory but an orchestra.

Imagine a world-class symphony orchestra preparing for a bold experiment. For decades, the ensemble has mastered the great classical works—Beethoven, Mahler, Stravinsky—pieces that demand discipline, precision, and predictability. Every musician knows their part. Every movement has been rehearsed to the millisecond. The conductor controls the dynamics, the phrasing, and the tempo. Nothing is left to chance.

Now, the conductor announces a new program. The first half will be traditional. But the second half will be something entirely different: a structured improvisation. The orchestra will begin with a recognizable theme, then the musicians will depart from the sheet music, listening to one another, responding to each other's ideas, and shifting melodies and rhythms in real time.

It is a thrilling idea—and utterly unnerving.

A violinist wonders, *How will I know when to follow and when to lead?*

A cellist asks, *What if the horns move in a direction I don't expect?*

A timpanist thinks, *What if I misread a subtle gesture from the conductor?*

The conductor, too, must adapt. Instead of controlling every note, they must guide the overall arc—shaping, nudging, interpreting the emergent music without constraining it.

When the performance finally begins, the improvisation is electric. New melodies emerge. Unexpected textures bloom. Possibilities open that no one

could have planned.

But, of course, there are also moments of tension. A wind player misreads a harmonic hint and heads down a path no one else follows. A percussionist senses a shift that the brass section misses. There are moments when the ensemble briefly loses unity—not disastrously, but enough to remind everyone how delicate the improvisation really is.

The orchestra discovers what every modern organization eventually learns: Precision is easy when the world is predictable. The real test is how you behave when the sheet music ends. And three things determine whether the music transcends or collapses.

First, the musicians must learn from each other *in the moment*, not after the performance concludes. Waiting for next week's rehearsal to adjust is too slow. The correction must happen now, as the music unfolds.

Second, the musicians must be hearing the same evolving theme. Without shared context—without a common understanding of the tonal center, the rhythmic pulse, the emotional direction—each musician plays brilliantly yet incoherently.

Third, there must be governance in motion. The conductor cannot dictate every bar, yet there must be a frame: boundaries of tempo, mood, continuity, transition. Governance becomes a living guide, not a static rulebook.

This is exactly what organizations must master in the Intelligence Age.

Agents are the improvising musicians; leaders are the conductors; context is the shared musical theme; and governance is the invisible structure that keeps the performance coherent.

Rule 8 concerns allowing agents to learn in the loop, not after the fact. Traditional organizations treat learning as a retrospective exercise. They analyze failures after the quarter ends. They update models once sufficient drift has accumulated. They re-examine processes after a costly exception has already occurred.

This is the equivalent of musicians saying, “Let’s review what went wrong at next week’s rehearsal.” But the music is happening now.

In the Intelligence Age, agents must learn from live signals. When the world shifts, they must shift with it. When a human corrects an agent’s action, that correction becomes part of the agent’s knowledge. When anomalies appear, they become training data immediately, not months later. Learning happens in motion, not in hindsight.

Rule 9 is about how humans and machines must act with the same context. Improvisation collapses if musicians are not listening to the same music. Similarly, agentic collaboration breaks down when humans and machines operate from different understandings of the business.

The machine might have perfect access to structured data but lack understanding of conversational nuances, the implicit knowledge in relationships, or the strategic importance of exceptions. Humans may possess all that nuance but lack visibility into the agent’s assumptions.

When contexts diverge, trust diminishes; when trust diminishes, conflicts increase; when conflicts increase, agents fail to scale. Shared context restores unity. It provides humans with confidence that agents understand what matters and equips agents with the information they need to act responsibly in human environments.

Finally, Rule 10 concerns the governance of coagency in motion, not the data at rest.

For centuries, musical scores have governed orchestras. But in improvisation, the governance must shift. It is less about dictating every note and more about shaping the emergence. The conductor moves from being a ruler to being a steward. So too must organizations evolve.

Traditional governance focuses on data at rest—security policies, access rules, retention standards. Important, but insufficient. When autonomous agents make decisions on the fly, governance must shift to the agents’ behavior:

their boundaries, escalation paths, feedback loops, alignment with values, and real-time oversight.

Governance becomes dynamic. It must anticipate drift, correct misalignment, and adapt as agents learn. This is governance not of artifacts but of *actions*.

The Third Pillar Links Promise to Performance

Together, these three rules form the operating discipline of an agentic enterprise. Rule 8 ensures that the organization's intelligence stays alive rather than calcifying. Rule 9 ensures that humans and machines act as one ensemble rather than as disconnected performers. Rule 10 ensures that this ensemble remains coherent, aligned, safe, and strategically guided.

The first pillar asked, "What must intelligence look like?" and the second asked, "How do we build and fund it?" The third asks, "How do we make it work to win?"

Execution is improvisation at scale. The organizations that win will be those that can learn as they go, coordinate in real time, and govern autonomy without stifling it.

This is the heart of the third pillar: *Intelligence becomes action only when the organization becomes a symphony—capable not merely of playing the music it knows but of creating new music in the moment.*



Strategic Intent

Invest to Monetize

Implement to Win



RULE 8

LET AGENTS LEARN IN THE LOOP, NOT AFTER THE FACT

Leadership and learning are indispensable to each other.

—JOHN F. KENNEDY

On paper, Kodak and Sears should have owned the future.

Kodak literally invented the digital camera. In 1975, a twenty-four-year-old engineer named Steve Sasson walked into a conference room in Rochester, New York, with a toaster-sized prototype that could capture a digital image, store it on a cassette, and display it on a screen. It was slow, clunky, and monochrome—but it worked. When executives asked him when this would catch up to the quality of film, he replied, “Fifteen to twenty years,” and he was mostly right.

It filed the patents—and buried the idea. The entire economic engine of Kodak revolved around film and chemical processing. A world of instant digital images didn’t just threaten a product; it threatened the corporate psyche. Kodak did continue to tinker with digital sensors and even launched early digital cameras, but emotionally, it remained a film company. Metrics,

incentives, and mental models all pointed back to silver halide and paper. The company *learned about* digital; it did not *learn with* digital. Sasson had invented the digital camera too early for an analog company such as Eastman Kodak to capitalize on.⁵⁹

Sears tells a similar story. For decades, it was the Amazon of the Industrial Age: the Sears catalog, national logistics, private labels, stored customer accounts—it all added up to a simple but powerful data and fulfillment system. When the web arrived, Sears had the customers, inventory, distribution centers, and brand advantage. What it lacked was a learning loop. Instead of turning its catalog into a lively digital marketplace, it doubled down on fragmented divisions and quarterly store metrics.⁶⁰ Competitors such as Amazon, with no legacy stores to protect, built systems that learned from every click and shipment. Sears optimized yesterday's business logic; Amazon birthed the business model for tomorrow. We know which survived.

Both companies exemplify the main risk Rule 8 highlights: Organizations that view data as a rearview mirror rather than a nervous system are vulnerable. They use information to reflect on the past, not to guide future decisions. In an agentic world, that is deadly.

From Reports to Reflexes

In the Industrial Age, data behaved like a report card. You closed the books each quarter, looked back at performance, and perhaps adjusted plans for the next cycle. Feedback arrived in thick binders or static dashboards long after action had been taken. That lag—between what happened and what you learned—was measured in weeks or months.

In the Intelligence Age, data is more like a bloodstream. Signals flow from the edges of the business back to the center and out again, continuously. Digital agents—pricing engines, routing algorithms, personalization models, scheduling systems—sense what is happening, act on it, and then adjust based

on what comes next. The loop closes in seconds or milliseconds, not in the next off-site planning.

Rule 8 is about that loop: Let agents learn in the loop, not after the fact.

The agent here is not just a chatbot. It is any system that can observe, decide, and act—whether it's an AI model routing a driver, a recommendation engine curating content, or a predictive maintenance agent scheduling repair. And crucially, these digital agents are paired with human agents—planners, operators, clinicians, creatives—who provide context, judgment, and purpose.

When the loop works, every interaction improves the system for the next one. When it doesn't, you end up like Kodak and Sears: awash in information but structurally unable to learn fast enough to matter.

The distinction is not subtle. Traditional BI asks, "What happened last quarter, and what should we do differently next time?" Agentic intelligence asks, "What is happening right now, what action should we take, and how do we know if it worked?"

The former is periodic. The latter is continuous. The former treats learning as an occasional event. The latter treats learning as the mechanism of operation itself.

To see this in practice, let's examine organizations where learning in the loop is already the difference between winning and merely surviving.

UBER LEARNS WHILE YOU RIDE

Think about what happens when you open the Uber app.

You tap a button; a car appears on the map; a driver is dispatched; a route is chosen; a price is set. To the rider, this feels almost trivial now. Under the hood, it's a dense web of digital agents making thousands of microdecisions: which driver to match, which route to favor, what price to show, whether to offer a shared ride, how to adjust incentives if demand is spiking across town.

Uber runs all of this on Michelangelo, its internal machine learning platform. Michelangelo manages the entire machine learning life cycle—data ingestion, feature management, training, deployment, and monitoring—for hundreds of models and thousands of experiments at any given time. Those models don't sit on shelves; they are continuously fed with real-time telemetry from trips: GPS traces, traffic conditions, driver accept/decline patterns, cancellation rates, payment friction, and more.

Every ride becomes a learning opportunity. If an ETA prediction consistently misses the mark in a specific neighborhood, the monitoring layer detects the error and initiates retraining. When a new incentive scheme leads to longer idle times for drivers, the system updates the utilization metrics accordingly. If a change in the app interface causes riders in one city to abandon the booking process more often, product teams notice this pattern almost immediately.

This is agentic intelligence at scale and speed: Digital agents propose and execute actions in real time. Human agents—city managers, ops teams, data scientists—monitor outliers, define new objectives, and push back when the models behave badly. The loop of signal → decision → outcome → adjustment never stops.

Crucially, Uber has learned that humans are not just supervisors of machines; they are colearners. Safety teams, for example, review rides flagged by AI as unusual—such as sudden hard braking, odd detours, or long stops. Sometimes the model may give incorrect recommendations, but the pattern might have an innocent cause. Each human judgment feeds back into the training data, helping to refine what *suspicious* means. Over time, the system becomes both more accurate and less intrusive. Performance, in this model, is not something you check at the end of the quarter. It is something you prove in motion, trip after trip.

AIRBUS'S SKYWISE: A SHARED BRAIN FOR THE SKIES

Aviation is the opposite of agile. Aircraft are certified for decades, routes and timetables are planned far in advance, and safety regulations are set and enforced slowly and carefully. Yet behind the scenes, aviation is becoming one of the most sophisticated examples of learning in the loop.

Airbus's Skywise platform is an open, cloud-based data environment that connects airlines, suppliers, and Airbus itself. More than 12,300 aircraft are now connected to Skywise, forming a shared data backbone that integrates telemetry, maintenance logs, sensor readings, and operational events. Airlines contribute their own fleet data; in return, they get predictive analytics that help them anticipate failures, optimize maintenance, and improve operations.

Here, each aircraft effectively becomes a digital agent. When a vibration pattern on an A320 engine over Singapore indicates the early stages of a specific fault, it is compared against thousands of other incidents in the Skywise graph. If the system detects a risk, it generates alerts not only for that plane and that airline but for similar aircraft flying elsewhere. A maintenance planner in Toronto can act before a failure strands passengers on the tarmac; a mechanic in Frankfurt can swap parts proactively rather than reactively.

Airbus and its airline partners have reported double-digit reductions in unscheduled maintenance and millions in life-cycle savings per aircraft by using predictive maintenance based on Skywise insights. But the deeper story is organizational.

For Skywise to work, airlines had to change their behavior. Maintenance teams accustomed to treating data as local had to accept a shared system. Engineers had to learn to trust alerts surfacing from patterns they could not see directly. Executives had to overcome the instinct to hoard operational data as proprietary and instead participate in a governance framework where data is anonymized, access is controlled, and benefits are mutual.

The learning loop here spans multiple layers: Aircraft sensors feed continuous data; models detect patterns, propose maintenance actions, and learn from outcomes; human engineers validate those actions, refine business

rules, and flag edge cases; and the insights flow back into design, operations, and safety standards.

Skywise is not just about making individual planes smarter; it is about turning the entire ecosystem into a learning organism. This is learning in the loop crossing organizational boundaries—a rule we'll return to.

AMAZON: A WAREHOUSE THAT TEACHES ITSELF

If Uber is about learning on the street, and Airbus and Netflix are about learning in the sky and on the screen, Amazon is about learning in the warehouse.

Amazon has now deployed more than a million robots in its fulfillment network and is rolling out new foundation models such as DeepFleet to coordinate fleets of mobile robots across its warehouses. Each robot is both an actuator and a sensor. As shelves move, orders are picked, and packages flow, they generate a vast stream of data on routes, congestion points, picking accuracy, and timing.

That stream feeds into simulation and optimization systems. When robots consistently slow down in a particular aisle at certain times of the day, the system identifies a bottleneck. Engineers run digital twin simulations—changing layouts, re-sequencing tasks, adjusting pick paths—to test improvements virtually. Once a new configuration is deployed, the robots' behavior provides immediate feedback: Did throughput rise, errors fall, worker ergonomic metrics improve?

This is not purely a robotic story. Human associates remain central. They are the ones who notice when a theoretically optimal layout causes unexpected strain, or when a new workflow confuses seasonal workers. Their observations, often captured in structured feedback tools or frontline apps, provide the contextual layer that pure telemetry cannot. In Amazon's case, agentic intelligence is a choreography: robots, algorithms, planners, and frontline workers learning from one another in a continuous loop.

The result is not just efficiency—it's adaptation. When demand patterns shift (e.g., holiday peaks, regional preferences, supply chain disruptions), the system doesn't wait for a strategic review before responding. It learns and adjusts, often before human planners even notice the pattern.

CARMAX: TURNING OMNICHANNEL INTO A LEARNING ENGINE

CarMax, the used-car retailer, is not an obvious example of agentic intelligence. But it illustrates how learning loops can transform even traditional businesses.

CarMax has turned its omnichannel platform into an engine of continuous learning—using unified data about customers, vehicles, and financing to adapt buying journeys, streamline operations, and experiment with generative AI summaries to help buyers decide faster. When a customer browses inventory online, schedules a test drive, negotiates financing, and completes paperwork—whether in-store, online, or in a hybrid format—the entire journey generates signals.

Those signals feed agents that optimize inventory placement (which cars to stock where), pricing (dynamic adjustments based on local demand and competitive positioning), and customer experience (what information to surface when, what offers to make). The agents don't operate in isolation—they're continuously refined by sales associates who understand local market nuances, finance teams that spot patterns in approval rates, and operations teams that manage vehicle transfers across locations.

This is learning in the loop applied to a business model that predates the internet. The key is that CarMax didn't just add AI to its existing processes. It rebuilt processes around the learning loop: Sense customer behavior → propose optimizations → execute changes → measure outcomes → adjust models. We at Reltio are proud to have helped CarMax gain an edge on this journey.

The Architecture of Learning: How to Build Your Own Loops

How does a company that has grown up on quarterly reports pivot to learning in the loop?

Most organizations begin with a familiar pattern: a big data lake, a BI team, and dashboards that summarize outcomes. Those assets remain useful, but they are fundamentally after the fact. They help you understand yesterday. To support agentic intelligence, the architecture must move closer to the operational edge.

Step 1: Instrument the key flows. In practice, this means instrumenting the key flows in your business—orders, claims, treatments, trades, shipments, interactions—with event streams. Each event becomes a small story: what happened, to whom, where, under what circumstances. Those events feed into a shared graph that captures entities (customers, assets, locations, contracts) and relationships (who bought what, which device is in which facility, which component belongs to which asset, which clinician treated which patient). This requires moving beyond batch extract, transform, and load thinking. Events must be captured when they happen, not consolidated overnight. Latency matters—minutes or seconds, not days.

Step 2: Deploy agents that observe, decide, and act. On top of that event stream and graph, you deploy agents. A fraud detection agent watches transaction events and proposes holds or alerts. A dynamic inventory agent watches sales, shipments, and supplier signals and proposes reorder quantities. A clinician support agent watches admitted patients, lab results, and risk scores and proposes care-path adjustments.

The key is that each agent's action is logged, monitored, and evaluated. Did the fraud hold prevent loss or annoy a valuable customer? Did the inventory adjustment reduce stock-outs or create new bottlenecks upstream? Did the care recommendation reduce complications or produce unintended side effects?

Step 3: Close the loop with evaluation and feedback. Those evaluations—ideally supported by well-designed experiments—are then fed back into the models. This is where many organizations stumble. They deploy the agent, see it work, and move on. But without systematic evaluation, you don't know *why* it worked, *when* it fails, or *how* to improve it.

Closing the loop requires (a) telemetry: capturing not just outcomes but intermediate states, confidence scores, and alternatives considered; (b) experimentation: running A/B tests or counterfactual simulations to isolate causal effects; (c) human-in-the-loop feedback: structured ways for operators to flag edge cases, propose corrections, or validate agent reasoning; and (d) retraining pipelines: automated or semi-automated processes to incorporate new data and push updated models to production.

This is socio-technical by design. Product managers and domain experts decide which outcomes matter and what trade-offs are acceptable. Data scientists decide which signals to include, and how to measure “better.” Operational leaders decide when to trust the agent fully, when to require a human in the loop, and when to block certain actions altogether. Governance teams ensure that learning does not come at the expense of safety, fairness, or compliance.

Step 4: Scale across the organization. Once a single loop is working—say, fraud detection or inventory optimization—the pattern can spread. The same architectural principles (event streams, shared graphs, agents with evaluation loops) apply across the supply chain, customer service, HR, finance, and operations.

The challenge is not technical replication but organizational change. Different departments must agree on common semantics for entities and events. They must accept that their local rules may be overridden by agents that learn from broader patterns. They must shift performance metrics from “Did we execute the plan?” to “Did the system learn and improve?”

This is where leadership matters. If executives treat each learning loop as an isolated AI project, they miss the compounding effect. If they recognize that learning loops are the new operating model, they create incentives and structures to support them.

Learning with the World, Not Just Within Your Walls

One more crucial dimension: Feedback loops increasingly cross organizational boundaries.

Airbus showed this in aviation, where Skywise creates shared learning across airlines, suppliers, and the manufacturer. Spotify shows it in music, where listener behavior is cocreated by labels, artists, and platforms. Amazon shows it in fulfillment, where robotics vendors, Amazon Web Services teams, and logistics partners all contribute to and learn from shared models.

In an agentic ecosystem, your systems will increasingly learn not just from your own data but from signals that originate in partner networks. A pharmaceutical supply chain might learn from anonymized hospital outcome data; a payments network might learn from merchant fraud patterns; an industrial OEM might learn from how operators actually use equipment in the field.

Rule 8, in other words, depends on the earlier rules: Data must be linked across boundaries (Rule 3) and trusted as a currency of coagency (Rule 4), or the loops collapse under suspicion. You can't have effective learning if partners fear their data will be weaponized against them or leaked to competitors. Trust infrastructure—technical, legal, and cultural—becomes the foundation for cross-boundary learning loops.

This is where competitive dynamics shift. In the Industrial Age, companies guarded operational data as proprietary. In the Intelligence Age, the winners will be those who can participate in learning ecosystems while protecting what

truly differentiates them. The trick is knowing which data to share (to benefit from network learning effects) and which capabilities to keep proprietary (to maintain competitive advantage).

Learning Loops Are Where Strategy Meets Reality

At first glance, Rule 8 might seem technically interesting but strategically secondary—something for the data team to address. We think it'd be a big mistake to feel that way.

Feedback loops are where strategy meets reality.

If Rule 1 (neural architecture) and Rule 2 (data in motion) describe how you build the nervous system, and Rule 3 (data ecosystems) and Rule 4 (trusted data) describe how you connect it beyond your own skin, then *Rule 8 is about how that system behaves over time*. Does it merely execute yesterday's plans faster? Or does it learn, adapt, and improve with every cycle?

Business executives—not data scientists—own that question. You decide whether a model's output is treated as a one-off project deliverable or as a living capability that must be monitored, tuned, and governed. You decide whether streams from the edge are wired into agentic loops or parked in lakes for someone to analyze later. You decide whether teams are rewarded for defending last year's logic or for experimenting with new loops that might disrupt comfortable habits.

The Kodak and Sears stories are not about “dumb” leaders ignoring “smart” engineers. They are about organizations that lacked the structures—and the will—to learn from their own inventions. Sasson's digital camera sat outside any meaningful loop in Kodak's business. Sears's early web efforts were bolted onto a legacy of siloed divisions. Both companies treated new data as an add-on, not as the center of gravity.

The companies we've examined—Uber, Airbus, Netflix, Spotify, Amazon, CarMax—are far from perfect. Many face regulatory scrutiny, labor tensions, or business model risks. But they share one discipline: They have hardwired learning into the way their agents operate. They treat intelligence not as a slide in a strategy deck but as a property of the system.

That is the heart of Rule 8.

The Leadership Mandate

First, stop treating AI as a series of projects, prototypes, or experiments. Projects have beginnings and ends. Learning loops do not. When you commission a model, you're not buying a finished product—you're starting a relationship that requires continuous care, feeding, and governance. Budget accordingly. Staff accordingly. Measure accordingly.

Second, demand visibility into the loops themselves. Don't just ask, "What did the model predict?" Ask, "How did the model learn from what happened next?" If your teams can't show you the feedback mechanism, you don't have agentic intelligence—you have glorified batch processing.

Third, align incentives with learning, not just execution. Reward teams that improve their agents' performance over time, not just teams that deploy models. Celebrate failures that generate useful feedback. Penalize gaming behaviors where people optimize for the metric while breaking the underlying loop.

Fourth, invest in the socio-technical infrastructure of learning. This means event architectures that capture actions and outcomes in real time; experimentation platforms that isolate causal effects; feedback tools that let humans correct agent behavior; governance frameworks that define when agents can act autonomously and when they need human approval; and training programs that teach people how to supervise, refine, and collaborate with learning agents (connecting back to Rule 7 on agentic literacy).

Finally, recognize that learning loops change the nature of strategy itself. In a world of quarterly planning cycles, strategy meant setting a direction and executing against it. In a world of continuous learning loops, strategy means defining objectives and constraints, then letting the system discover the best path.

You're not commanding and controlling. You're teaching and steering.

Bringing It Back to Agentic Intelligence

Agentic intelligence, as we've defined and discussed throughout this book, is the collaborative capacity of humans and machines to act coherently in a complex environment. It is not just smarter models but smarter organizations.

Allowing agents to learn in the loop makes that collaboration tangible. This is how companies gain an edge.

When a digital agent proposes an action, sees the outcome, and adjusts; when a human agent can see what the system saw, test alternatives, and feed judgment back into the model—then you no longer have AI projects. You have a living, learning enterprise.

In that world, quarterly reports still matter—but as narrative summaries of what your loops have already learned, not as the primary means of discovering what's going on. Strategic reviews still matter—but they are informed by a thousand small experiments already run, not guesses about a world you haven't tested.

The companies that master this will find that each decision makes the next decision easier, not harder. They will accumulate what you might call *learning capital*: a compounding advantage embedded in their agents, their data, and their people.

This is the essence of execution in the Intelligence Age. Not just acting faster but learning faster. Not just deploying smarter technology but building organizations that get smarter with every cycle.

The Risk of Loops Without Governance

One final note of caution: Rule 8 is powerful, but it becomes hazardous if learning loops are misaligned or ungoverned.


A loop that optimizes for the wrong objective can learn its way into disaster remarkably quickly. A fraud detection agent that minimizes false negatives might block so many legitimate transactions that customers flee. A content recommendation engine that maximizes engagement might amplify outrage and misinformation. A hiring agent that learns from historical patterns might encode and amplify existing biases.

This is why Rule 8 cannot stand alone. It depends on (a) shared context (Rule 7): Humans must understand what agents are learning and why; (b) trusted data (Rule 4): The signals feeding the loop must be accurate and representative; and (c) governance (Rule 10, which you'll see later): clear frameworks for defining objectives, monitoring behavior, and intervening when loops go wrong.

But the central point stands: *If you want agentic intelligence, you cannot bolt learning on at the end. You must put it in the loop.*

The organizations that understand this—that build learning into the fundamental architecture of how they sense, decide, and act—will discover something remarkable: The more they operate, the smarter they become. Not because they hired smarter people or bought better technology, but because they designed a system that learns.

That is the crux of Rule 8. That is the path from execution to evolution.



RULE 9

HUMANS AND MACHINES MUST ACT WITH THE SAME CONTEXT

It is amazing what you can accomplish if you do not care who gets the credit.

—HARRY S. TRUMAN

On paper, both systems looked smart.

In practice, they failed for the same reason: humans and machines weren't working from the same picture of reality.

A few years ago, Amazon experimented with an AI résumé screener designed to speed up hiring. The idea was simple enough: Train a model on ten years of past résumés and hiring decisions, then let it rank new applicants from one to five stars. But the data feeding the model reflected a male-dominated tech workforce. The result: The system quietly downgraded résumés that included the word *women's* (as in “women's chess club”), penalized graduates of women's colleges, and favored language patterns common in male candidates.⁶¹

On its own terms, the model was working—it had learned to mimic Amazon's historical hiring patterns. But the recruiters using it were operating

with a different mental model: They believed they were searching for the best candidates, not re-encoding historical bias. The AI and the humans were acting from *two very different contexts*. Once that became clear, Amazon rightfully scrapped the tool.

Boeing's 737 MAX crisis was the mirror image of the same problem. To address a new aerodynamic quirk in the aircraft, engineers introduced a software layer called the Maneuvering Characteristics Augmentation System (MCAS) that could automatically lower the nose of the aircraft if it detected a stall. Pilots, however, were barely told it existed. When a faulty sensor fed MCAS incorrect data during two separate flights, the system repeatedly forced the nose down. Pilots fought back on the controls, but they were wrestling an invisible adversary they didn't know they had. Two planes crashed, 346 people died, and investigators around the world concluded that unintended interactions between human crews and hidden automation were central to the tragedy.⁶²

In one case, the humans treated the AI as a neutral assistant, even though it amplified old biases. In the other, the software silently overruled the humans who were ultimately responsible for safety. Both failures share a common root: The humans and the machines never shared a clear, aligned view of the situation they were acting in. They were not looking at the same map. The companies may have had exemplary systems of record; one of them (Amazon) is considered world-class for its systems of engagement with customer data and purchase graphs that drive superior recommendations. What they lacked was a shared system of context for the two types of agents to see the world. This is the execution failure Rule 9 addresses.

You can invest in the best AI models (Rule 5), build agentic literacy across your workforce (Rule 7), and design learning loops (Rule 8)—but if humans and machines don't share the same operational context in the moment of action, execution breaks down. Value is unrealized.

Rule 9 is about implementation: ensuring that when humans and machines execute together—approving a loan, routing a truck, diagnosing a patient, scheduling a flight—both are acting from the same understanding of what is happening, why it matters, and what “good” looks like right now. This isn’t strategic intent. It isn’t an investment. It’s how you deliver value to stakeholders as coagents to win in the marketplace. This is how you improve customer satisfaction levels and enhance your net promoter score.

Without shared context, AI becomes either a dangerous black box or a glorified spreadsheet: something you can’t trust or something you can safely ignore. With shared context, coagency becomes an operational lever in the Intelligence Age.

What We Truly Mean by Context—And Why It’s a Problem in Implementation

This book is organized around three pillars: Strategic Intent (Rules 1–4), Invest to Monetize (Rules 5–7), and Implement to Win (Rules 8–10). We developed Rule 9 to sit squarely in the third pillar because shared context can be designed, but it is useful only when it drives implementation.

You can have perfect strategic intent—a vision of human–AI collaboration and a commitment to digital transformation with shared context in designing the protocols and decision rights. You can make all the right investments—models, platforms, and training programs. But when a nurse stands at a bedside deciding whether to escalate care, when a dispatcher chooses whether to override a routing algorithm, or when a loan officer reviews a credit decision—that’s the moment of truth for how well a system is designed for agents to reason and act. If humans and machines don’t see the same situation at that moment, implementation fails.

In traditional IT, we discuss “single sources of truth”—a clean customer master, a reconciled ledger, a canonical product catalog. These are important,

but they don't tell the whole story. Context is more than just a table of records. It's the dynamic picture of what's happening right now, why it matters, and what constraints exist. At Reltio, we highlight the importance of context and promote the role of the system of context within our clients.

Think about driving. The raw data might be sensor readings: distance to the car ahead, speed, GPS coordinates, and lane markings. Context is the understanding that "I'm merging onto a busy highway in the rain, the car to my left is drifting, and the truck behind me is accelerating." That's the kind of situation model every competent human driver carries in their head.

Modern AI systems are, increasingly, built to hold situation models too. Autonomous driving teams discuss "scene understanding." LLMs operate within "context windows"—slices of documents, emails, chats, and past interactions that they can attend to at once. AI practitioners have begun to describe this design discipline as "context engineering": the craft of deciding what information an agent should see, in what structure, at what moment, so that its actions make sense in the real world.

But notice: If only the machine sees that context, you haven't solved the problem. You've just moved it. For humans and machines to act as coagents, they must share enough of that situational model that they can understand and critique each other's actions. That's the only way they amplify each other's strengths.

That's what this rule is about. Not just data quality. Not just dashboards. A shared frame of reference that makes collaboration possible.

Three Dimensions of Shared Context

For humans and machines to execute together effectively, you must implement shared understanding across three layers. These aren't theoretical categories—they're operational requirements that determine whether coagency works in practice.

1. SHARED PERCEPTION: BOTH MUST SEE THE SAME SITUATION IN REAL TIME

At the moment of execution, both parties must be looking at the same facts. When a credit agent declines a loan, the loan officer must see the same credit history, the same risk factors, the same customer signals. When a routing algorithm suggests a delivery path, the driver must see the same traffic conditions, delivery windows, and constraints the algorithm is optimizing for.

This seems obvious, but it's routinely violated in implementation. Agents train on one dataset while humans make decisions using another. Models update overnight while human playbooks and routines remain frozen. The machine sees real-time signals, while the human relies on yesterday's report or what they learned in last year's training programs or corporate retreats.

Implementation question: Are humans seeing the same situational data as your agents when they act? If the answer is yes, you have created the right system of context.

2. SHARED INTERPRETATION: BOTH MUST UNDERSTAND WHAT IT MEANS

Raw facts aren't enough for execution. A sepsis alert means nothing to a nurse if they can't see *why* the model flagged it—rising heart rate, falling blood pressure, changes in lactate. A fraud hold frustrates a customer service rep if they can't determine *the pattern* that triggered it—transaction velocity, device mismatch, or geographic anomaly.

Shared interpretation means both parties understand not just *what* happened, but *why it matters* for the action they're about to take. The machine must make its reasoning visible. The human must understand how the machine weighs evidence.

Implementation question: When your agents recommend an action, can humans see the reasoning behind it in language they understand? If the answer

is yes, you have the right system of context for agentic workflows.

3. SHARED OBJECTIVES: BOTH MUST OPTIMIZE FOR THE SAME OUTCOMES

Even when humans and machines see the same situation and understand the same patterns, execution breaks down if they're optimizing for different goals. A recommendation engine that maximizes engagement online might amplify outrage, while the editorial team is optimizing for trust. A scheduling algorithm that minimizes cost may create unsustainable workloads for employees, while the operations manager is optimizing for retention.

Shared objectives mean both parties are working toward the same definition of *good*—and both understand what trade-offs are acceptable when carrying out their tasks.

Implementation question: Are your agents optimizing for the same outcomes your humans care about in their operations, and are they aligned with corporate priorities and culture? If the answer is yes, you are well on your way to realizing consistent value from deploying agents at scale.

When all three layers align in execution—perception, interpretation, and objectives—humans and machines can genuinely act as partners. When one—even one—layer fractures, coagency breaks down at the moment it matters most: at the point of action.

From Dashboards to Glass Cockpits: Aviation's Lesson

Pilots have a useful phrase for this: staying in the loop. When commercial aviation first introduced sophisticated autopilots, there was a wave of accidents that didn't stem from mechanical failure but from mode confusion. Pilots didn't always know which systems were active, what logic they were following,

or how the autopilot would behave next. They were, quite literally, out of the loop.

Modern cockpits are designed very differently. Displays emphasize a mental model of the flight—where the plane is, what’s controlling it, and what’s likely to happen in the next few minutes. Automation still flies most of the time, but pilots see what it sees. They can predict its actions, override them, and understand why it took those actions. *The cockpit is a shared context surface.*

We’re now replicating that same shift across every industry that adopts agentic systems. The question is no longer “Should we automate?” but “How do we design systems where humans and machines share the same situational understanding in daily execution?” This is the core reason why agentic literacy (which builds on data literacy) is crucial, and we highlighted it as a key area of investment in Rule 7.

How Leading Companies Implement Shared Context

The companies winning with agentic intelligence are not those with the most advanced models. They are the ones that have embedded shared context into day-to-day operations—so humans and machines act on the same reality, at the same time.

Microsoft Copilot is often described as a productivity tool, but its real strength lies elsewhere. Copilot is grounded in Microsoft Graph—a living map of emails, meetings, documents, chats, and organizational relationships. When a manager asks for a summary of customer conversations, Copilot doesn’t generate text in isolation; it traverses the same work context the manager inhabits. Crucially, the manager can correct it—exclude internal debates, add a missed call, refine relevance. Those corrections reshape the shared

representation. Productivity emerges not from automation, but from coediting context.

On the factory floor, Toyota applies the same principle in physical form. At its Indiana plant, Toyota uses edge AI systems from Invisible AI to monitor manual assembly steps. The system flags anomalies, but humans decide what they mean. When a team lead confirms a real defect—or dismisses a false alarm—that judgment feeds back into both the AI model and the human standard of work. Workers and machines learn together, side by side. The AI behaves less like a supervisor and more like a junior apprentice, improving through shared experience rather than detached observation.

In logistics, UPS's ORION system shows how shared context scales across thousands of daily decisions. ORION proposes optimized routes based on global data. Drivers follow them—but also override them when local reality disagrees. Those overrides are not treated as exceptions; they are treated as signals. When humans consistently deviate from the model, the model changes. Over time, the map and the territory converge. Optimization improves precisely because humans and algorithms work from—and reshape—the same situational picture.

Finally, Airbus Skywise demonstrates the shared context across organizational boundaries. Airlines contribute telemetry and maintenance data to a governed platform, but they also see the collective interpretation: emerging fault patterns, comparable incidents, and effective interventions. A maintenance planner in Toronto and an engineer in Frankfurt act from the same contextual understanding. Skywise works not because data is centralized indiscriminately but because context is shared selectively, transparently, and with mutual benefit.

Across software, factories, logistics, and aviation, the pattern is consistent. Agentic intelligence does not emerge from better models alone. It emerges when humans and machines see the same situation, can question each other's conclusions, and learn from the same feedback loops. Shared context is the

quiet architecture that makes that possible—and without it, even the most powerful AI remains isolated, brittle, and slow.

Implementing Shared Context

So, what does all this mean for you as a leader trying to execute with agentic intelligence? This is not about strategy documents or investment committees. This is about how you run operations when humans and machines must act together.

1. MAKE CONTEXT DECISIONS AT EVERY SYSTEM DESIGN CHOICE

First, recognize that every architectural decision is also an implementation decision about shared context. You are making deliberate choices about what your agents see—which tables, which events, which documents, which signals from the outside world—and how that picture is exposed to humans at the moment of action.

Every decision about data architecture, system integration, and interface design determines whether humans and machines can execute together or will work at cross-purposes.

Ask yourself: When you approve a new system, do you ask, “Will humans and machines see the same situation when they need to act?”

2. DEMAND OPERATIONAL TRANSPARENCY, NOT JUST MODEL PERFORMANCE

Second, insist that agents make their reasoning visible at the point of execution. If your credit-risk model declines a loan, a frontline banker must be able to see the logic in language they understand *at that moment*. If an operations agent reallocates inventory from one region to another, planners

must see not just the new plan but the assumptions behind it *when they receive it*. If an HR matching engine suggests candidates for a role, recruiters must recognize the factors it implicitly weighs *as they review candidates*.

This isn't about dumbing down the AI. It's about making the AI's reasoning legible enough that humans can execute with confidence—or intelligently override when context demands it.

Ask yourself: Can your people explain why an agent recommended what it did, or are they just following instructions?

3. BUILD FEEDBACK LOOPS THAT CLOSE IN OPERATIONAL TIME

Third, build the plumbing so that human corrections flow back immediately. When a nurse overrides an alert, a driver rejects a route, or a product manager flags a recommendation as off base during their shift, those actions should flow back into the learning loop before the next shift, not disappear into a quarterly review.

This requires structured feedback mechanisms built into operational tools (not separate survey systems). It needs attribution systems that link decisions to outcomes in near real time. It needs retraining pipelines that incorporate human judgment within operational cycles (days, not months). And it surely needs governance processes that decide when to trust human corrections and when to investigate—executed by people close to operations.

Ask yourself: How long does it take for a frontline correction to update the agent's behavior?

4. ORGANIZE AROUND CONTEXT, NOT FUNCTIONS

Finally, organize operations so that humans and agents share the same view. Delta's Operations and Customer Center works because dispatchers, pilots, meteorologists, and AI systems all see the same integrated picture. Mayo's

digital twin works because clinicians and diagnostic agents both see the unified patient record. Singapore's Smart Nation works because city planners and algorithms both access the same sensor fabric.

This often means breaking down functional silos that create distinct operational realities across the organization.

Ask yourself: Do your operational teams work from the same context their agents use, or do they have separate systems?

Agentic Literacy Includes Context Literacy

In Rule 7, we argued that data literacy—the ability for employees at all levels to read, question, and use data—is now a basic corporate skill. It's necessary but not sufficient.

A data-literate employee can interpret a chart. A context-literate employee can understand how and why that chart fits into a wider pattern of behavior, incentives, and constraints—and how an AI agent will act if it sees that chart.

Likewise, an agent-literate organization doesn't just train people to click the right buttons in a tool. It teaches them to reason about how their digital coworkers see the world: What inputs are the agents using? How do those relate to frontline reality? When should we follow their suggestions, and when should we deliberately contradict them? What happens to the model when we do?

Tesla's evolution of FSD is a messy but instructive case. The company has built one of the world's largest human–AI learning loops by training on real-world driving data from millions of cars and using human driver interventions—disengagements, corrections, overrides—as key training signals. The entire point of end-to-end neural networks for driving is that they learn not just from labeled video clips in a lab, but also from the actual road context.

At its best, that's what agentic learning looks like: Machines practice in the real world, humans provide live feedback, and both gradually converge on safer

behavior. At its worst, as ongoing safety investigations remind us, it shows what happens when context isn't fully shared—when drivers misunderstand the system's limits, or regulators don't see the same data the company does.

The lesson for other industries is not to copy Tesla's controversies but to heed the underlying point: *Performance and trust are inseparable from shared context*. If employees and regulators are left guessing about what an agent knows, coagency breaks down. Agentic intelligence isn't relevant.

Why Shared Context Wins

When humans and machines execute with the same context, organizations gain three decisive operational advantages:

- **Faster execution cycles.** Decisions happen at system speed, not meeting speed. Teams don't stop to reconcile conflicting information. Models don't wait for humans to provide missing context. Everyone—every agent—executes from the same operational picture.
- **Higher precision in action.** Ambiguity disappears at the point of execution. Whether approving a loan, rerouting a delivery truck, adjusting a flight schedule, or issuing a price quote, actions reflect the shared current state of the world—not competing interpretations.
- **Coherent multiagent execution.** The organization executes like a coordinated system, not a collection of conflicting departments. Competitors still reconcile siloed views; you execute with a unified understanding.

This is the foundation of strategic advantage in the age of agentic intelligence. Not better strategy. Not more investment. Better execution through shared operational context.

What Happens When Context Splits

The stories of Amazon's biased résumé screener and Boeing's MCAS failures weren't just about flawed models. They were about broken context:

- The AI at Amazon saw a different reality than human recruiters.
- The pilots of the 737 MAX were flying with a different operational context than the machine that controlled the aircraft.

Split context = split judgment = catastrophic outcomes.

In business terms: A recommendation engine that optimizes for one objective while the sales team operates under another creates customer friction; a supply chain model forecasting high demand while procurement sees shortages generates bullwhip effects; a fraud algorithm flagging suspicious behavior that customer service agents cannot interpret damages trust.

The cost isn't merely inefficiency—it's competitive decay. Organizations with fractured context move slowly, waste capital, frustrate customers, and lose their ability to coordinate at scale.

Meanwhile, competitors with unified context operate like well-oiled machines (excuse our mixed metaphor from the Industrial Age).

Think Context—Not Prompt—Engineering

Andrej Karpathy of OpenAI and Tesla fame has argued that modern AI systems are only as robust as the context window they operate within—the structured, relevant environment they're given to act upon. We strongly believe that the same principle applies to companies as they navigate toward the Intelligence Age.

Context engineering is the emerging corporate ability to design environments where humans and machines interpret the world consistently. It involves unified real-time data, not scattered logs; shared semantic layers, so

terms mean the same thing across roles and models; aligned objectives, so agents optimize for the same outcomes that humans value; transparent reasoning, so humans can understand why an agent suggests an action; and rapid correction loops, so machine errors are quickly reinterpreted through human judgment.

The aim is simple: Ensure every agent—human or digital—is operating from the same understanding of the situation. Companies that master this will outperform those that do not, just as companies that mastered integrated supply chains outperformed those stuck in siloed operations.

Shared Context Is How You Get an Edge

Winning in the Intelligence Age doesn't come from having more AI. It doesn't come from better strategy or bigger investments.

It comes from *better implementation through intelligence that's aligned between humans and machines.*

The companies that will dominate the next decade are those where execution happens seamlessly because: Humans trust the machine's recommendations because they understand the operational context behind them; machines incorporate human corrections because they're fed back into the model in operational time; and both execute inside the same world—the same data, the same signals, the same goals, the same real-time feedback loops.

That's not automation; it's coordinated implementation. That's how you win.

When humans and machines execute within the same context, agency stops being a tug-of-war. The machine isn't a mysterious oracle; it's a partner whose reasoning you can inspect, challenge, and improve during operations. The human isn't a slow, error-prone bottleneck; they're a source of rich, situational signal the system can't get anywhere else—*in the moment of action.*

That is the essence of symbiotic intelligence: two very different kinds of minds inhabiting the same operational map, executing toward the same goals.

Our message is this: Shared context is not a technical choice. It is an execution imperative. And it may be the most important implementation decision a company makes in the age of agentic intelligence.

In the next and final rule, we'll address the hardest execution challenge: Once you have agents learning in the loop and acting with shared context, how do you govern that agency in motion—in real time—without grinding execution to a halt?

[OceanofPDF.com](https://oceanofpdf.com)



RULE 10

GOVERN AGENCY IN MOTION, NOT DATA AT REST

You can't steer something you can't control.

—MUSTAFA SULEYMAN, CEO, MICROSOFT AI

Control. As agentic intelligence progresses, it becomes increasingly important—because we're creating autonomous neural networks of data that act on their own to speed up and protect our processes.

But we have to remember that there's this thing called Murphy's Law.

Murphy's Law states that anything that can go wrong will go wrong. In the Intelligence Age, it's an apt warning, because unanticipated interactions can wreak havoc if control is not established.

The world found that out the hard way on July 19, 2024.

A routine update issued by CrowdStrike propagated automatically across millions of machines. Within minutes, hospital terminals froze, airline check-in systems went dark, banks locked users out, retailers shut down point-of-sale networks, and ERs reverted to paper and walkie-talkies. The disruption

spanned continents and sectors, unfolding faster than any human response could keep pace with.⁶³

What made the event so unsettling was not the scale alone. It was the cause. There was no cyberattack, no malicious actor, no breach of perimeter defenses. A software agent did exactly what it was designed to do: Update itself. And in doing so, it revealed something far more consequential than a technical flaw.

It revealed that we had built systems capable of operating at a global scale, at machine speed, without designing a way to govern their actions in motion.

That moment is now burned into the executive consciousness because it exposed something more fundamental than a technical flaw. The failure wasn't one of data governance. The data was secure. The pipelines worked. The update passed internal tests.

The failure was one of *agentic governance*—the governance of systems that sense, decide, and act on their own while the world scrambles to keep up.

This is what defines the final rule—that now is the time to refocus governance on our agentic intelligence in motion, rather than stored, static data.

The Library Versus the Highway

Governing static data is like managing a library. You organize the collection, control who has access, keep records of what was checked out. Everything sits on shelves. You can audit at your leisure. If someone misfiles a book, you'll find it eventually.

In contrast, governing AI agency is like managing highway traffic. Cars are moving at seventy miles per hour, making split-second decisions about lane changes and braking. You can't audit after the accident. You need traffic lights, speed limits, lane markings—constraints that operate at the speed of motion itself. Without them, you don't just get a messy library. You get a pileup. Either

you have guardrails, signals, and speed limits operating in real time, or you have chaos moving at seventy miles per hour.

However, as we've seen over and over again throughout these pages, old habits die hard. After all, for decades, governance was built around a simpler premise: Govern the data, and the systems will behave. Data was stored, cataloged, secured, and audited. Systems were passive. Humans were the locus of decision-making. Governance could afford to be procedural, periodic, and external because action itself was slow. Oversight lagged execution, but that lag was tolerable.

That world no longer exists.

Today's systems aren't waiting for instructions. They learn continuously from streaming data, update parameters dynamically, coordinate with other systems, and act in real time. They decide which transactions to block, which routes to reroute, which customers to engage, and which risks to escalate—often millions of times per hour. In such an environment, governance that exists outside the system or arrives after the fact is no longer governance at all. It becomes a postmortem analysis.

Rule 10 follows naturally from everything that came before. Once data is linked rather than siloed, once it flows in motion rather than resting in reports, once humans and machines act from shared context, and once agents learn in operational loops, a new question becomes unavoidable:

Who constrains agency while it is executing?

The CrowdStrike outage made that question impossible to ignore. It showed what happens when autonomous action outpaces the mechanisms designed to supervise it. There was no circuit breaker for agency, no way to pause propagation while understanding what was unfolding. The world did not lack intelligence at that moment. It lacked restraint at speed.

Restraint at Speed

To see what effective restraint looks like, it helps to study an industry that has been governing agency in motion for decades. Commercial aviation offers the clearest archetype. For more than half a century, aviation has coordinated pilots, autopilots, air traffic controllers, maintenance systems, and regulators while operating at hundreds of miles per hour. It does so not by slowing execution but by embedding governance directly into it.

Aviation's Safety Management System was never designed as a compliance exercise. It is an execution framework. Authority is distributed but bounded. Pilots can override automation, divert flights, or declare emergencies, but always within well-defined limits and under continuous observation. Automation is powerful but transparent. Pilots do not interact with black boxes; they see modes, constraints, projected paths, and underlying logic. They can anticipate the system's next action and intervene when context requires it.

Equally important, the system learns from execution. Every flight generates data not just about outcomes but about deviations, overrides, and near misses. These signals flow into global learning systems that update procedures, training, and safeguards. Governance itself evolves through use. The result is not rigidity but resilience—an industry capable of acting at speed without sacrificing safety.

In contrast, Zillow's iBuying failure illustrates what can go wrong when restraint is not a part of the planning process when it comes to agentic intelligence. In 2018, the real estate company thought it had found a way to use AI to revolutionize home buying. Its algorithm would predict future home prices, buy properties directly from sellers, make minor improvements, and flip them for profit. They called it "Zillow Offers."

They then gave the algorithm the authority to make purchasing decisions worth hundreds of thousands of dollars per transaction. Per its programming, it analyzed comparable sales, neighborhood trends, market conditions, and recommended which houses to buy and at what price.

For a while, it seemed to work.

Until it didn't.

By late 2021, Zillow owned thousands of homes—and the algorithm had dramatically overpaid for most of them. Home price appreciation slowed. Renovation costs exceeded projections. The algorithm kept buying at inflated prices even as market conditions shifted. The program was shut down that November. The damage? Eight hundred eighty-one million dollars in losses and the loss of 25 percent of their workforce because of layoffs.

From a governance perspective, Zillow made several critical missteps that ultimately contributed to its algorithmic failures. The company granted the algorithm too much autonomous authority, allowing it to commit to home purchases worth hundreds of thousands of dollars with minimal human oversight—prioritizing volume over accuracy. It also failed to implement adequate circuit breakers, lacking automated mechanisms to pause operations when the algorithm's predictions diverged significantly from market conditions. Compounding the issue was insufficient real-time validation; Zillow did not monitor closely enough whether projected prices aligned with actual resale values in time to mitigate financial losses. The algorithm operated as a black box, making decisions that even company executives struggled to interpret, which made it difficult to intervene when things went wrong. Finally, the system lacked graceful degradation. It was designed for speed and scale, but when failures occurred, they were catastrophic rather than incremental, leaving no room for recovery before significant damage was done.

The algorithm had learned from historical data when home prices were rising steadily. When market conditions changed, it couldn't adapt—and the governance structure didn't catch the problem until hundreds of millions of dollars were already lost.⁶⁴

The Question of Authority

For the first time in organizational history, we are delegating authority—not just execution—to systems that operate continuously, invisibly, and at scale. These systems do not simply recommend. They decide. They act. And their actions propagate instantly across customers, markets, and institutions.

Whenever societies have crossed a threshold like this, governance has followed disruption, not preceded it. Financial markets learned this after repeated crashes, when unbounded leverage outpaced oversight. Industrial systems learned it after catastrophic failures, when safety standards emerged only after tragedy. Early aviation learned it through accidents that forced the creation of air traffic control, maintenance regimes, and shared safety protocols. In every case, the pattern was the same: Capability surged first; governance arrived later, under pressure.

The Intelligence Age is unfolding on the same trajectory—but at a much faster rate.

Agentic systems are already embedded in payment rails, logistics networks, healthcare workflows, credit decisions, energy grids, and information ecosystems. They are shaping outcomes before humans are aware a choice has been made. Yet many organizations still treat governance as something to be layered on once systems are stable, or delegated to compliance teams once value has been proven.

The truth is, we can't wait until Murphy's Law kicks in. We have to put in place governance early in the process.

When agents act at machine speed, the absence of governance is not neutral—it is a decision. It is a decision to allow systems to optimize locally without regard for systemic effects. It is a decision to discover failure only after customers, citizens, or markets experience harm. And it is a decision to trade short-term velocity for long-term legitimacy.

This is why governing agency in motion is not simply an operational discipline; it is an institutional one. It determines whether intelligent systems earn trust or erode it, whether autonomy compounds value or amplifies

fragility. In a world where systems act continuously, legitimacy is not conferred by intention or reputation. It is earned through behavior—visible, bounded, and correctable behavior—at the moment of action.

The uncomfortable truth is that no regulator, auditor, or external authority can solve this for organizations. Governance at the speed of action must be designed by those who build and deploy agentic systems. It must be embedded in architecture, encoded in rules, surfaced through transparency, and reinforced through learning loops. Anything less is symbolic control.

Future historians of technology will likely look back on this period not as the moment AI became powerful but as the moment organizations decided whether they were willing to be accountable for that power. The winners will not be those that moved fastest, but those that learned how to move fast without losing control.

And that responsibility now sits squarely with leadership.

Across these seemingly disparate domains, a common pattern emerges. Authority is never absolute; it is bounded. Constraints are not documented and hoped for; they are embedded and enforced. Oversight is continuous rather than episodic. Learning is systemic, not incidental. Governance happens at the same speed as execution, not after it.

This is the heart of Rule 10. Governing agents in motion requires a fundamental shift in how organizations think about control. Controls must live where agents act, not where data is stored. Autonomy must be defined explicitly, with clear escalation paths when boundaries are reached. Agency must be observable, so humans can see not just what agents are doing but why. And governance itself must learn, incorporating overrides, failures, and exceptions into ever-better constraints.

When these conditions are absent, organizations face a false choice. Either they rein in autonomy, slowing execution until governance can catch up, or they unleash autonomy and accept fragility at scale. Neither is sustainable.

Competitive advantage in the Intelligence Age belongs to those that can move fast *and* remain in control.

This is why governance is no longer a cost center or a compliance function. It is execution infrastructure. It determines how confidently an organization can deploy agents, how quickly it can adapt, and how much trust it can sustain as systems scale. The companies that win will not be those with the smartest models but those that have designed architectures where intelligent action unfolds within visible, enforceable, and evolving bounds.

Ultimately, the Intelligence Age is not defined by whether machines can act autonomously. That question has already been answered. It is defined by whether humans are willing—and able—to design systems where autonomy operates responsibly, transparently, and in alignment with intent.

Strategy at the speed of data demands governance at the speed of action. Without it, intelligence becomes risk. With it, agency becomes advantage.

That is why this rule comes last. And why it matters most.

What Makes It Agentic Governance?

At this point, you could reasonably ask: Isn't this just continuous improvement with fancier tech?

Not quite. Traditional continuous improvement methods—Lean, Six Sigma, Kaizen—were powerful, but they were primarily human-driven, periodic, and localized. Teams measured, tested, and tweaked, but the pace and scope were limited by human throughput.

Agentic governance introduces three crucial differences.

First, it operates at machine timescales. Uber's surge pricing doesn't wait for a weekly review; it adjusts minute by minute. Amazon's routing agents don't replan warehouses annually; they learn daily. Netflix doesn't revise its home page once a quarter; it experiments continuously. The feedback cycle

compresses from months to milliseconds, allowing thousands or millions of micro-adjustments that accumulate into a strategic advantage.

Second, it is embedded in shared datagraphs, not isolated spreadsheets. Netflix's view of audience engagement is linked to content metadata, device data, and geographic behavior within one integrated system. Airbus's outlook on engine health is based on a graph of flight conditions, component histories, and fleet-wide patterns. This reflects your earlier rules: architecture as a neural network (Rule 1), data in motion (Rule 2), cross-boundary ecosystems (Rule 3), and trusted data as coagency (Rule 4). Without that foundation, you get noise, not learning.

Third, and most importantly, it formalizes coagency. Digital agents are not just passive calculators; they take actions that affect customers, assets, and employees. Human agents are not just button-pushers; they shape objectives, interpret anomalies, and decide when to override or retrain. In a well-designed loop, each side teaches the other. The machine proposes; the human disposes; and the outcome becomes training data for the next iteration.

That reciprocity is the essence of Rule 10. Without it, you have what might be called "batch intelligence"—models trained once, deployed statically, and evaluated occasionally. With it, you have genuine agentic intelligence: systems that adapt with every decision, informed by both machine pattern recognition and human judgment.

How the Three Final Rules Connect

Rules 8, 9, and 10 create an implementation triad. They're not separate principles—they're interconnected capabilities that must work together for organizations to succeed with agentic intelligence.

Rule 8 emphasizes that agents must adapt to real-world conditions rather than rely solely on initial training data. Learning is an ongoing process embedded within workflows. However, learning without proper governance

leads to chaos. If agents change their behavior based on every signal without limits, they may stray from strategic goals. Learning from biased feedback can also cause errors to escalate. Rule 10 determines what agents learn, how quickly they adapt, and when human oversight is needed. Governance doesn't halt learning; it ensures that learning aligns with organizational goals.

Rule 9 guarantees that execution occurs consistently. Both humans and machines perceive the same situation, understand the same reasoning, and aim for the same goals. However, shared context without proper governance can lead to alignment on the wrong objectives. If humans and machines both see biased data, they will make biased decisions together. If they share outdated information, they will act in sync but based on old assumptions. Rule 10 oversees the quality of shared context—ensuring it is accurate, timely, and properly scoped for the decisions at hand.

Rule 10 ensures agents operate within boundaries, escalate appropriately, and act transparently. However, governance without learning becomes fragile. Static rules cannot keep up with changing environments. Governance must evolve as agents learn. And governance without shared context becomes ineffective. If governance systems monitor agents using signals that differ from those operators use, efficiency declines when humans must intervene.

The three rules together create a tight system: Agents learn continuously; humans and machines execute from a shared understanding; governance ensures that learning and execution remain aligned with strategic intent. This is how organizations implement to win in the Intelligence Age.

The Strategic Edge of Governance

Let's return to the competitive question: Why does agentic governance determine who wins?

Because in a world where every company has access to similar AI models, the differentiator isn't the intelligence—it's the implementation.

India's Unified Payments Interface (UPI) thrives not because it is a better payment technology than that of other countries, but because it can execute 640 million daily transactions with governance embedded at transaction speed.⁶⁵ Competitors that can't govern at this speed either throttle growth to maintain control or scale chaos, leading to fraud and failures.

Aviation wins not because planes are foolproof, but because the industry operates billions of passenger miles with disciplined, controlled agency. Industries that audit safety quarterly rather than monitor it continuously can't achieve the same safety record.

When governance becomes part of how you execute—not something that happens after execution—you gain three decisive advantages.

First, faster execution cycles. You don't slow down to ensure safety. Safety is built into speed. UPI completes transactions in under two seconds, including fraud checks, authentication, and logging. The governance doesn't add latency—it is the execution.

Second, scalable trust. You can execute across organizational boundaries and millions of autonomous decisions, as is done with UPI, because governance creates trust at scale. Trust isn't built through contracts and handshakes. It's built through infrastructure that makes violations impossible or immediately visible.

Third, continuous improvement through execution. Your governance system learns from every execution like aviation's Safety Management System does. You don't just prevent yesterday's failures—you anticipate tomorrow's risks. The governance architecture becomes smarter with every transaction, every override, every exception.

Organizations that master governed execution don't just avoid catastrophic failures such as CrowdStrike's outage or Zillow's iBuying debacle. They execute faster, learn quicker, scale further, and capture more value than competitors that still treat governance as a compliance exercise. Governance is no longer a

defensive capability. It is the execution capability that determines competitive position.

Govern Agency in Motion, or Be Governed by the Consequences

The Intelligence Age leaves every organization with a simple, unavoidable choice: Will you govern the intelligence you deploy in execution—or be governed by the consequences of failing to do so?

The answer depends on decisions you make now. Where do humans stay in the loop during execution? When do machines act autonomously? How does context flow across agents in real time? What signals trigger escalation? What must be explainable at the point of action? What must be reversible? How do you encode ethics into operational architecture? How do you ensure trust survives scale?

These are not policy questions. They're implementation questions. And the answers determine whether your organization can execute in the Intelligence Age or merely survive in it.

We began this book with a simple idea: The Intelligence Age changes more than technology. It changes the structure of competition, the role of leaders, the meaning of data, the shape of organizations, and the relationship between humans and machines.

The ten rules were not written as a rigid framework or a checklist. They were written as a field guide—a way to navigate a transition as profound as electrification, industrialization, and digitization before it. But among all the rules, this one is the hinge. The one that determines whether the others can execute at all.

The first rule showed that data must be linked like a neural network—but a network is only reliable if governance ensures the pathways are trustworthy.

The second rule emphasized velocity—but speed without guardrails is amplified liability, not advantage. The third rule pushed organizations to cross boundaries—but interconnection exposes firms to cascading risks if governance is inconsistent. The fourth rule elevated trusted data as the foundation—but trust dissolves the moment machines act on stale or unauthorized inputs.

The fifth, sixth, and seventh rules showed how to invest across horizons, monetize action, and prepare people—but investments don't scale unless governance provides a stable architecture. The eighth rule taught that agents must learn in the loop—but continuous learning is only safe when governance provides oversight, correction, and accountability. The ninth rule insisted that humans and machines must share context—but shared context collapses without shared governance ensuring visibility, traceability, and alignment.

And Rule 10 completes the system. It is the switch that turns strategy into execution, investment into results, and capability into competitive advantage.

You do not govern data. You govern decisions. You govern actions. You govern agents. In doing so, you govern how your organization executes in the Intelligence Age.

And if you do it well, you won't just adapt to this era—you'll execute your way to the front of it. This is how organizations turn agentic intelligence into a competitive reality.

EPILOGUE

To close, let's review the story of a fictional movie and television production company, EverLight Studios.

EverLight had always believed it understood enchantment and entertainment. For nearly a century, its stories had filled theaters, living rooms, theme parks, and imaginations worldwide. Parents introduced children to the films they had loved as kids. Families made pilgrimages to its parks, much as earlier generations once saved for grand European tours. Toy aisles bent around its franchises. Streaming platforms measured themselves against its catalogs. Even rivals conceded—quietly—that EverLight didn't merely entertain. It shaped culture.

And then, almost imperceptibly at first, the light began to dim. It didn't start with a crisis. It started with confusion.

The First Cracks

A highly anticipated animated series launched on EverLight's streaming platform to fanfare and record viewership projections. But within hours, customer service lines lit up. Viewers in Tokyo, São Paulo, and Berlin reported placeholder subtitles scrolling across their screens. Dialogue tracks were out of sync by many seconds—an eternity to the viewers. In some markets, characters spoke lines that contradicted the visuals—scenes that had been recut months ago but somehow persisted in certain language versions.

Inside EverLight's gleaming headquarters, the finger-pointing began immediately.

The localization team in Dublin swore they had delivered the proper versions—on time, properly tagged, exactly to specification. The platform

engineering team in Los Angeles insisted they had ingested the correct files and could produce logs proving it. The legal affairs group in London said the file names didn't align with the rights documentation, meaning certain versions should never have been released. The New York marketing team said no one had informed them of any localization issues, and they'd already committed to the global launch date months ago.

Everyone was right. Everyone was wrong. And no one could see the whole picture.

Inside the executive suite, people chalked it up to growing pains. EverLight had always run complex operations. Entertainment on a global scale was messy. Mistakes happen. Engineering would fix the subtitles, corporate communications would issue a press note, and everyone would move on.

Except the mistakes kept happening.

A Pattern Emerges

Three weeks later, a consumer products team in Hong Kong released a toy line centered on a character arc the creative team in Burbank had rewritten six months earlier. The toys hit shelves showing the character as a villain—but in the final film, she'd become the hero. Confused parents flooded social media. Retailers demanded explanations. Marketing scrambled to spin the discrepancy as “expanded universe storytelling,” but the damage was done.

Then, international broadcast partners in Germany and France received master files containing scenes that had been removed from the domestic theatrical release—scenes that violated content standards in those markets. The broadcasts were pulled mid-premiere. Regulatory fines followed. Partnership agreements were renegotiated at a high cost.

More trouble. EverLight's streaming algorithm began recommending darker, mature-audience content to family accounts because metadata tags had never been harmonized between the theatrical, home video, and streaming

systems. Parents complained. Regulators opened inquiries. Subscriber growth stalled.

Even more trouble. A European licensing partner recalled an entire product line after discovering that EverLight no longer held rights to the character depicted—a change in ownership that had occurred during a corporate acquisition two years earlier. The legal team knew. The licensing team did not. The contract management system had never been updated. The recall cost tens of millions.

None of these failures was fatal. But together, they formed a pattern—one that leaders could no longer ignore.

EverLight, a company built on stories, had lost its own data—and business—story.

The World Had Moved On

In the opening pages of this book, we sketched out AirMed—a hospital that transformed itself into an intelligent organization, where every piece of data connected to every other, where humans and machines worked from shared context, where learning happened in real time, and where governance operated at the speed of action. AirMed's early warning signs were clinical: slowly rising wait times, inconsistent patient records, the iconic shoebox, tiny cracks in a system that hinted at deeper architectural gaps. The leadership stepped up to address them.

EverLight's warnings were subtler still, buried in a world that had moved faster than the studio understood.

While EverLight struggled to coordinate thirty different metadata profiles for a single film across fifty disconnected teams, competitors were releasing localized content simultaneously across dozens of markets with no friction. Their recommendation algorithms improved daily because production, marketing, and platform data lived in the same digital network. Their

merchandising teams worked from unified metadata graphs that connected characters to products to audiences in real time. Their stories flowed effortlessly between screens, formats, and geographies—theaters to streaming to airlines to theme parks—because every system spoke the same language.

One rival studio had already seen the future. Years before EverLight stumbled, that competitor unified its creative, legal, rights, localization, and distribution data into a single living ecosystem. Every character, every asset, every line of dialogue, every contractual clause, every localization nuance flowed through one shared architecture. When that studio launched a global franchise, every territory received the correct version, every merchandising partner had current character data, every marketing campaign reflected the final creative, and every licensing agreement was validated in real time.

The result was not just operational efficiency. It was creative freedom for the innovators—both human and digital.

Teams could experiment faster because they weren't drowning in coordination overhead. Content could move globally with confidence because everyone saw the same truth. AI models could enhance—not endanger—decisions because they operated on trusted, connected data. The studio didn't become less magical. It became more so.

That rival's operational excellence caught the attention of industry giants. Major studios and streamers began circling, recognizing that in the Intelligence Age, content libraries alone weren't enough. You needed the architecture to activate them. Partnership discussions turned into acquisition talks. The studio that built its nervous system first became the target of every major player. And at unprecedented price premiums.

EverLight had all the same resources. All the same talent. All the same creative legacy. It simply waited too long to build the foundation.

The Day Everything Collided

The breaking point came with a film intended to reintroduce EverLight to the world—a global tentpole with A-list talent, groundbreaking visual effects, a story designed to resonate across cultures, and a marketing campaign that spanned continents. It was the company’s boldest release in years, meant to prove that EverLight still commanded the magic that defined modern entertainment.

Marketing had spent eighteen months building anticipation. Retailers across forty countries had stocked shelves. Streaming partners had committed to premium home page placement on launch day. Airlines had been contracted for in-flight versions timed to theatrical windows. Theme parks in three continents had prepared character rollouts, interactive experiences, and merchandise launches for the same weekend.

Everything was carefully planned. Everyone was excited.

And then, six days before the global premiere, a regional legal team in Eastern Europe flagged a discrepancy in the rights documentation.

The character at the center of the film—EverLight’s newest icon, designed to anchor a decade of sequels, spin-offs, attractions, and consumer products—was not cleared for distribution in several major territories. Or perhaps it was. Or maybe the rights had been renewed during a corporate restructuring but not recorded in all systems. Or perhaps the contract had been updated in one repository but not propagated to others. Or perhaps a subsidiary acquisition had complicated ownership, but no one had traced the full chain of title.

No one could say for sure.

Legal teams in four countries began emergency reviews. Technology teams tried to reconstruct the audit trail across seventeen different contract management systems, three rights databases, and countless email chains. Creative executives insisted the character was original IP. Business affairs executives produced documents showing otherwise. Engineers searched for the single source of truth and discovered there wasn’t one.

The film slipped from its release date. A new date was announced. It slipped again. Then, it moved to the following quarter. The media didn't react kindly, and the stock market reaction was ruthless.

Millions in marketing spend evaporated. Retailers marked down pre-positioned merchandise at a loss. Streaming partners reallocated home page real estate to competitors. Airlines renegotiated contracts. Theme parks pushed character rollouts to the following season—if at all. Analysts openly questioned EverLight's operational competence. The stock price tumbled. Partnership discussions with potential acquirers cooled. Rivals took note of the weakness.

The problem wasn't the film. The story was brilliant. The production was flawless. The marketing campaign was inspired.

The problem was the business foundation—the logic of how to compete in the Intelligence Age—beneath it.

What EverLight Never Built

AirMed reimagined its entire organization as a nervous system—signals moving instantly, connecting staff, patients, machines, and contexts into a cohesive whole. Data didn't sit in silos. Intelligence didn't move in batch processes. People didn't work from fragments. Systems learned from one another. Governance operated in motion, not in retrospect.

EverLight had never done that work. It thought its storytelling magic was more than enough to keep it successful. After all, it had a time-honored brand, creativity, talent, heritage, and global reach. What it lacked ... was a digital nervous system.

Its data lived in isolated pockets, scattered across acquisitions, legacy platforms, and departmental fiefdoms. Its intelligence moved in slow motion, passing through email chains, meeting agendas, and quarterly reviews. Its people saw only fragments of the whole, and its systems didn't learn from one another; they simply coexisted. Its governance relied on human memory,

institutional knowledge, and heroic effort—not on architecture, automation, and real-time validation.

EverLight still believed it was competing in the entertainment sector using Industrial-Age logic. What the studio didn't realize was that it was now a player in a very different game, a game where real-time data within a shared context and intelligence rule—and they were losing.

Every entertainment company now competes on the same dimensions that define every other digital business: How fast can you learn? How well can you coordinate? How intelligently can you act? How effectively can you govern at scale? Content is necessary but not sufficient. Creativity is essential but not decisive. The winners are the organizations that can activate their creativity with intelligence—and the losers are those that treat data, systems, and architecture as back-office problems to be solved eventually. Sometime in the future but not this quarter or this fiscal year because the budgets are allocated to existing projects and priorities.

A Different Story Was Possible

The rival studio that caught the industry's attention—the one that Netflix, Amazon, Apple, Google, and Paramount all reportedly considered acquiring—didn't just build better content. It built a better infrastructure—one suited for the world of agentic intelligence.

It treated data as a strategic asset, not an IT problem. It integrated creative, legal, rights, localization, and distribution data into a dynamic graph in which every node informed every other. It was designed to ensure that when a creative decision changed, downstream systems updated automatically. It built governance into the architecture so that rights violations, localization errors, and data conflicts were flagged before content reached markets—not after.

It invested in agentic literacy across the organization so that creative teams understood how their decisions propagated through the system, and

technology teams understood what creative teams needed to express. It let agents learn in operational loops so that recommendation algorithms improved daily, localization quality rose continuously, and rights validation became proactive rather than reactive. It ensured humans and machines worked from a shared context so that when a marketer launched a campaign, it saw the same character data that merchandising, licensing, and platform teams used.

It governed agency in motion so that when content moved across borders, formats, and partners, every guardrail activated automatically—validation happened in real time, compliance was built into the workflow, and quality was enforced at the point of action.

That studio didn't become less creative. It became operationally unstoppable.

And when industry consolidation accelerated, when streaming wars intensified, when content libraries became acquisition targets, that studio became the prize everyone wanted because it had the one thing you can't buy or fake: an intelligent architecture that could activate its assets on a global scale.

EverLight could have built the same thing. It had the resources. It had the talent. It had the foresight to know change was coming.

It simply never treated the shift from the Industrial Age as an urgent priority. It fell into its own competency trap.

The Story That Comes Back to Us

Every leader reading this book has a choice between building an AirMed-style enterprise—a company that learns, synchronizes, adapts, and acts with intelligence—or drifting, slowly and almost invisibly, into EverLight's pattern of fragmentation.

EverLight didn't collapse overnight. No dramatic announcement. No single catastrophe. Just a thousand small moments where intelligence was possible—

and never built. A thousand decisions to defer the hard work of connecting systems, unifying data, investing in architecture, building governance, and treating intelligence as infrastructure rather than aspiration.

The future doesn't arrive as a crisis. It arrives as a slow accumulation of disconnections.

And then one day, you realize your competitors are moving faster, coordinating better, learning continuously, and governing intelligently—while you're still stitching together emails, reconciling spreadsheets, and hoping someone remembers how the pieces connect.

The Question That Remains

You've reached the end of this book. You've read ten rules. You've seen concepts, cases, and implementation patterns. You've met organizations that transformed themselves and understood what's at stake.

Now comes the only question that matters: Which story are you writing?

AirMed's story—the hospital that saw the transformation coming and built its nervous system before crisis forced the issue?

Or EverLight's—the studio that had every advantage, every resource, every reason to transform, but waited too long?

In the opening pages, we asked you to imagine AirMed's future—a morning in 2030 when intelligence was woven into every interaction, every decision, every moment of care. Systems that learned. Humans and machines that worked from a shared context. Governance that operated at the speed of action. Not because it was forced. Because it was designed.

Now, as you close this book, imagine *your* organization's future one morning—say Wednesday, February 20, 2030:

- What will your systems know that they don't know today?
- What will your human experts see that they can't see now?

- How will they collaborate with machines—whether in the cloud or in humanlike form?
- What will your organization be able to do that's impossible today?
- How quickly will you be able to learn, adapt, and act?

And most importantly: Will your architecture enable intelligence, or will you still be fighting your own infrastructure?

The Intelligence Age doesn't care about your industry, your legacy, or your past dominance. It rewards only one thing: the ability to learn faster, coordinate better, and execute more intelligently than your competitors.

AirMed revealed what happens when leaders design for intelligence early. EverLight highlighted what happens when they don't.

The ten rules in this book are not predictions, nor prescriptions, nor aspirations for every company. They're rules that are likely to shape the architecture of survival and success in an era when every company is becoming an intelligence company—whether it realizes it or not.

You can build your nervous system now, while you have the resources and the runway. Or you can build it later, after competitors have moved, customers have left, and partners have found alternatives.

But you will build it. You will be forced to do it.

The only question is whether you'll do it from a position of strength—or necessity.

In a world where every company is becoming an intelligence company, what story is your organization writing right now?

And when others tell that story years from now, will it sound more like AirMed's—or EverLight's?

The choice, as it always has been, is yours.

But the opportunity to choose is running out. We hope you will make a good choice and decide to write the script now.

AGENTIC TRANSFORMATION AS THE NEXT ARC OF STRATEGY

BY VENKAT VENKATRAMAN

When I think back on forty years of studying technological change—from the mainframe era to the rise of the internet, the explosion of mobile computing, the platform revolution, and now the emergence of autonomous agents—one pattern is clear: Each generation tends to misinterpret the next. We look ahead based on what we already understand. We assume things will stay the same when big changes are coming. We tell ourselves that the next wave will act much like the last one, maybe faster or stronger, but still similar. We see new technology through old perspectives—and in doing so, we miss the full change before it happens.

I've seen this pattern repeat across different industries and over many decades. And honestly, I've fallen into the same trap myself.

The Assumption I Didn't Know I Was Making

In the early 1990s, while I was a faculty member at MIT Sloan School of Management, I was working on MIT's Management in the 1990s program. At that time, the internet had not yet become mainstream, and the idea that companies would fundamentally reorganize around digital networks seemed speculative at best—radical and impractical to many observers. However, through that research, I began to notice an underlying logic that would

eventually shape not just one transformation but three decades of continuous change.

In 1994, I published an article in *MIT Sloan Management Review*⁶⁶ explaining how IT would go beyond automating existing tasks to redefine the firm's core boundaries and scope. That prediction, which felt bold then, proved surprisingly enduring. The logic I described back then became the foundation for understanding the rise of ERP systems, CRM platforms, e-commerce, mobile computing, cloud infrastructure, and ultimately the platform economy that now supports much of global commerce.

That framework influenced my later work—*The Digital Matrix*,⁶⁷ *Fusion Strategy*,⁶⁸ and several academic articles—and each version seemed to push the boundaries of what technology could do within organizations. But behind all these frameworks, no matter how progressive they appeared at the time, was a core assumption that felt so natural, so obvious, that I hardly noticed I was making it: *Humans would remain the main decision-makers*. Technology would enhance our abilities, improve our judgment, and boost our intelligence—yet it would fundamentally aim to make *us* more effective. The focus would stay on human cognition and human choice.

That assumption, which structured decades of thinking about digital transformation, is now dissolving before my—and our collective—eyes.

The Shift That Changes Everything

What we're witnessing now is not just the next phase of digital transformation; it's the emergence of something fundamentally different. We're entering a time when intelligence is no longer confined to human minds or fixed automation rules. Instead, intelligence *arises* from systems—networks and groups of humans and machines that collaborate, learn from each other's strengths, and continuously adjust to new conditions. In this environment, software does more than automate routine tasks—it actively collaborates with human

judgment. Workflows don't just follow predefined steps—they learn from results and adapt their behavior accordingly. Organizations don't just optimize within set boundaries—they evolve instantly, guided by agents capable of reasoning through complex issues, negotiating trade-offs, and taking actions that would have seemed like science fiction just ten years ago.

This is why I see agentic transformation not as just the next chapter in the digital story we've been telling for thirty years, but as an entirely new narrative arc. Digital transformation mainly asks: *How can we make this process faster, cheaper, or more accurate? How do we eliminate waste, minimize errors, and scale operations?* These questions focused on optimization within current frameworks. Agentic transformation asks a more fundamental question: *Who—or what—should be doing this work? What becomes possible when decision-making isn't limited by human capacity or machine inflexibility but instead results from the collaboration between complementary forms of intelligence?* The difference isn't just one of degree—it's a fundamental shift. Digital transformation changed how companies operate, improving the mechanics of execution. Agentic transformation changes how companies think, transforming the core nature of organizational intelligence itself.

What Four Decades Have Taught Me

Throughout this book, we've articulated ten rules that together form what I think of as the conceptual grammar of agentic transformation. These rules span diverse contexts—healthcare systems grappling with patient safety, media companies managing global content distribution, financial institutions balancing risk and speed, logistics networks coordinating millions of shipments, manufacturers integrating supply chains, government agencies serving citizens—yet beneath this variety runs a consistent set of principles. Intelligence emerges from *architecture*, not from individual applications operating in isolation. Learning becomes *continuous* rather than episodic,

woven into the fabric of daily operations rather than confined to periodic training cycles. Value gets created through *coordination* across boundaries rather than through isolated optimization of individual processes. Agency becomes *shared* between humans and machines in a genuine partnership, not ceded entirely by one to the other. And perhaps most fundamentally, trust must be *engineered* into systems through transparency, explainability, and governance—it cannot simply be assumed or mandated through policy.

These principles didn't arise from theoretical ideas alone. They are based on forty years of research on how organizations succeed or fail when technological change forces major strategic shifts. I observed this pattern with ERP systems in the 1990s, when companies had to choose whether to update their processes to fit integrated software or stick with fragmented legacy systems. I saw it again with e-commerce in the 2000s, when established retailers had to decide whether to treat online as a separate channel or to rethink their entire customer experience. The pattern repeated during the platform revolution of the 2010s, which changed global competition by emphasizing orchestration over ownership, and once more with the cloud–AI merger of the early 2020s, which brought computational intelligence at scales and speeds that were previously impossible. In each case, the organizations that succeeded weren't necessarily those with the most resources or the longest histories—they were the ones whose leaders looked beyond tools and technologies and toward fundamental changes in how value is created, captured, and maintained.

Agentic transformation refers to a recurring pattern that is amplified to an unprecedented level. It challenges leaders not only to automate tasks currently performed by humans but also to fundamentally *change* the nature of work itself—shifting from simply deploying systems to *designing relationships* between humans and digital colleagues that learn, adapt, and act with greater sophistication.

Why This Moment Feels Different

I've spent forty years studying how technological change reshapes organizations, and I've learned to be cautious about calling any particular shift unprecedented or transformational—those words are often used too loosely in business talk. But I believe that agentic transformation is the most significant organizational shift I've seen in my career, and I don't say that because the underlying technology is more advanced than what came before (even though it is), or because the pace of change is faster (though it definitely is). I say it because this transformation fundamentally changes what organizations are capable of *becoming*.

Digital transformation, despite its significant impact, mainly provided us with better tools to do what we already knew how to do. We could process transactions faster, reach customers more effectively, coordinate operations more smoothly, and analyze patterns more comprehensively. These were substantial advances that transformed industries and generated substantial value. But agentic transformation offers something fundamentally different: It grants organizations new *abilities* that weren't possible before. The capacity to learn at machine speeds while incorporating human judgment. The ability to coordinate across organizational, technical, and geographic boundaries that once seemed unchangeable. The capacity to adapt instantly to conditions that couldn't have been predicted or programmed ahead of time. The power to act with a form of distributed intelligence that grows and improves, rather than reaching a limit at the capacity of individual human cognition or automated rules.

The organizations that master this shift won't simply be more efficient or more profitable versions of what they are today—they'll be fundamentally different *kinds* of organizations, capable of operating in ways that would be genuinely impossible under either purely human agency or purely automated systems. That's what makes this moment feel, to me, categorically different from the transformations I've studied before.

Two Stories, One Choice

The story of AirMed that opens this book shows what becomes possible when organizational intelligence is built from the start—when leaders recognize early on that the future will demand new ways of coordination, learning, and decision-making. Conversely, the story of EverLight that ends the book illustrates what happens when that foundational work is postponed, delayed, or dismissed as something that can be addressed eventually—when leaders believe they can continue operating with fragmented systems, isolated data, and sequential workflows even as competitors develop integrated systems that learn and adapt in real time.

Although these two companies are fictional, they are based on ideas I've observed while studying both successful and less successful organizations during past waves of transformations. These patterns have repeatedly appeared across various industries, regions, and time periods over the last four decades. During technological shifts, the winners are rarely the largest organizations, those with the longest histories, or those with the deepest pockets. Instead, they are those that build highly adaptable *systems*—ones capable of sensing environmental changes, learning from outcomes, coordinating responses across different functions, and executing quickly and precisely through shared understanding between humans and machines. Success isn't determined by accumulating the most data in huge repositories but by turning data into coordinated intelligent actions faster than competitors can even recognize what's happening. The winners are not necessarily those deploying the most advanced AI models alone but those who have built the foundational *architecture*—the integrated databases, governance structures, learning loops, and shared context mechanisms—that enable humans and machines to collaborate effectively, improve continuously, and expand their capabilities over time, rather than reaching the limits of what any individual human or algorithm can achieve alone.

The Arc Ahead

Agentic transformation isn't a destination you reach after completing a set of initiatives or checking items off a transformation road map. It's not a project with an endpoint where you can declare victory and move on to the next strategic priority. It's something more fundamental and ongoing: It's a new way of leading organizations, a fresh approach to learning from experience, and a new framework for imagining what becomes possible when human judgment and machine intelligence work in genuine, continuous partnership rather than in the traditional relationship of human direction and machine execution.

The shift demands a fundamental redefinition of your organization. In an era where intelligence grows exponentially—where systems learn from every interaction, where coordination occurs on previously unimaginable scales, and where adaptation takes place in real time rather than through periodic strategic reviews—competitive advantage no longer primarily goes to those who merely adopt AI tools or implement machine learning models. Instead, it belongs to those who redesign their organizations as intelligent systems, where learning is continuous and integrated into daily operations, where context is genuinely shared between human and machine decision-makers, where governance moves at the speed of action rather than after the fact, and where humans and machines offer truly complementary strengths in every key decision. The organizations that will succeed in the coming decades won't be those that simply attach AI to existing structures and call it transformation. They will be those willing to do the harder work of reimagining their entire operating model around the opportunities that emerge when intelligence becomes distributed, collaborative, and ongoing.

The Question That Still Lingers

As you finish this book, having explored ten rules, various frameworks, and stories of both successful change and caution, one question ultimately matters

more than all the rest: Will you lead this transformation from a position of strength and choice, or will it happen to you under conditions of competitive pressure and organizational crisis?

The question isn't whether agentic transformation will reshape your industry—that's already happening, whether you've noticed it or not. The real question is whether you'll take the opportunity to develop your organization's nervous system while you still have the resources, the runway, and the strategic flexibility to do it effectively. Or will you find yourself trying to do that foundational work under pressure, rushing to respond after competitors have already acted, after customers have already changed their expectations, and after the market has already signaled that old methods are no longer enough? The chance to choose the former instead of being forced into the latter is available right now, but because of how quickly technological shifts happen, such windows don't stay open forever—they close, often faster than expected, and once they're gone, they don't reopen.

Where Theory Meets Practice

This book represents, in many ways, the culmination of a long intellectual journey. But it's also very intentionally the beginning of another conversation, one that takes place between strategy and practice, between academic research and executive experience, and between conceptual frameworks and the messy realities of implementation. That's where my collaboration with my coauthor becomes crucial to the book's purpose.

Manish Sood provides a practitioner's perspective on these frameworks—drawing on direct experience leading the development, deployment, and governance of autonomous agents at scale in competitive markets where mistakes can be costly and stakeholder patience for experimentation is limited. He explains how he has applied these rules at Reltio—as Customer Zero—and for his clients.

Together, our voices are meant to reflect the essential interaction that defines agentic transformation: Theory informs practice by providing clear concepts and recognizing patterns, while practice shapes theory by demonstrating what truly works in the real challenges of organizational life, competitive pressure, and human resistance to change.

OceanofPDF.com

RELTIO AS CUSTOMER ZERO

BY MANISH SOOD

When I launched Reltio in 2011, I was convinced of two things. First, that large enterprises were sitting on their most valuable asset and not getting much value from it: their own data. Second, the traditional way of managing that data was fundamentally broken for the world we were moving into.

More than a decade later, SAP's plan to acquire Reltio stands as strong validation that this problem was not peripheral, but central to the future of enterprise software and AI.

As I noted in this book's preface, I had spent years watching Fortune 500 companies make high-stakes decisions based on fragmented, inconsistent, and often outdated information. Every major platform claimed to be the system of record, yet customer, product, and supplier data were scattered across CRM, ERP, billing, marketing, data warehouses, and niche applications.

Very quickly, we realized that if we were serious about fixing this problem, we had to start with ourselves. Before we could ask global enterprises to run their business on Reltio, we had to become Customer Zero—using our own platform to run our business.

Data management tools were created to reduce fragmentation and give companies a reliable “source of truth.” But many of the older approaches were built for a different era: they were rigid, ran in scheduled batches, and assumed everything lived on company-owned infrastructure.

That mismatch showed up fast. Records were reconciled overnight while the business changed hour by hour. And the focus was often limited to a single

slice of the business, when companies increasingly needed a connected view across customers, products, suppliers, locations, accounts, and partners.

We believed there had to be a better way and started using a cloud-native approach that treats data as a living, connected asset and delivers trusted information at the speed of business, not the pace of yesterday's batch window.


Building the Reltio Data Cloud

Reltio was built around a simple idea: if the business runs in real time, your data has to keep up. That's why we designed our cloud platform from the start to be shared securely across customers, easy to connect to through APIs, and able to update information as events happen, not hours or days later.

We also aimed higher than creating a single "best" record for a person or product. We connect the dots between key entities and what's happening around them over time, including relationships, interactions, and transactions. That's what we mean by a 360-degree view: a complete, connected profile that pulls together what you know, what you've done, and how everything is related, all in one place and kept up to date.

For example, a customer 360 might link an account's parent company and subsidiaries to the right contacts, locations, contracts, recent orders, open invoices, and support tickets. So when sales, support, finance, and AI tools look at that customer, they're working from the same current picture and can act with confidence.

Reltio Context Intelligence Operating Model



Unify and mobilize your core enterprise data.

Over time, this evolved into what we now call the Reltio Data Cloud—a unification layer that integrates data from thousands of sources, enforces quality and governance, and makes trusted data accessible to any application, anywhere in the enterprise, in milliseconds.

As we worked with global pharmaceutical companies, retailers, insurers, financial institutions, and technology firms, they pushed us to define what a system of context needed to be. They didn't just want to know, Who is this customer? They wanted to know, What have they done? What are they likely to do next? Who are they connected to? What agreements govern the relationship? What happened last and why did it happen?

Answering those questions required more than unifying and reconciling core attributes. It required integrating structured data, behavioral signals, transactions, and, increasingly, unstructured data such as emails, notes, and call transcripts.

To support that reality, we invested in an ecosystem as much as a product. We support continuous integration of data from any enterprise data application or data sources and the delivery of the unified data context to any consuming application, business process, or analytics database. These contributing and consuming systems include:

- analytic platforms, such as Snowflake, BigQuery, and Databricks
- go-to-market systems, including Salesforce, Veeva, and HubSpot
- major ERPs and industry-specific applications across key sectors

This unified data context is further strengthened with support from integration of third-party enrichment from providers such as Dun & Bradstreet, Moody's Orbis, IQVIA, and ZoomInfo. We introduced industry-specific packages (which we refer to as "velocity packs") that bundle data models, connectors, and best practices commonly needed in industries, such as life sciences, consumer goods, and financial services.

The idea was straightforward: Our customers shouldn't have to begin with a blank slate. They already work within a complex environment. Our role is to unify data context in that complex environment to deliver simplification to support speed and scale.

Intelligent Data Is the Future

Over the past decade, our work has evolved from modern cloud data management into something bigger: the realization that clean, connected data is necessary, but not sufficient. As business leaders begin to adopt large language models and AI agents, they're also seeing a broader shift in how

software gets used. The “interface” is moving from screens full of menus and forms to chat boxes and task-focused agents that can take action on a user’s behalf.

That shift is already putting pressure on legacy SaaS. If users can ask an agent to “handle this workflow” instead of clicking through dozens of screens, the advantage moves from the user experience to the quality and availability of the underlying data. Many incumbent platforms see that threat, which is why some are tightening their grip on data inside their applications. But the Intelligence Age won’t work if data stays trapped in silos. Organizations will need data that’s available across the enterprise, continuously updated, and protected by strong security and governance.

That’s the emerging category we call *Context Intelligence*.

In that sense, SAP’s intention to acquire Reltio is not just a milestone for our company; it is broader recognition that enterprise AI will depend on trusted context spanning both SAP and non-SAP worlds.

The Intelligence Age won’t be unlocked by models alone. Models matter, but the real differentiator is an intelligent data layer that gives those models a deep, accurate, current understanding of the business, so agents can act safely and correctly. Context Intelligence is that layer. It brings together data unification, data quality, and structured and unstructured information, and presents it in a way that AI, analytics, and operations can use reliably.

A simple way to think about it is a city. The ontology is the map: the major things in the city and how they connect. The semantic layer is the rulebook: what each thing means, which sources are trusted, what policies apply, and what’s allowed.

Together, these two layers form a practical “system of context”:

Ontology layer: how you organize the business into clear concepts like entities, relationships, interactions, and groups. It’s what lets you model real-world complexity, like an employee working in a specific country under a specific retirement policy, or a healthcare provider, plan, and member

operating under contract terms and constraints. This structure helps AI interpret business questions as business logic, not just text.

Semantic layer: the meaning and guardrails around the data. This includes business definitions, metadata, lineage, policies, and classifications. It encodes how a metric is calculated, which system is authoritative, what regulations apply, and how data can and can't be used. This is what enables an AI agent to understand not just that a field exists, but what it means, when to trust it, and how to combine it responsibly with other information.

When you represent these layers as a graph, you get a powerful outcome: business meaning expressed as connected nodes and relationships that AI agents can navigate. That matters because agents increasingly need to answer questions that require multiple steps across real business context, like:

“Find all customers in this segment, in these markets, with this risk profile, whose usage has dropped in the last ninety days, and who are covered by contracts with this clause.”

Traditional flat data structures struggle with requests like that. A context graph, backed by a real-time operational data layer, is built for this scenario.

This is where Reltio's approach is different. The semantic graph explains what everything means. The Reltio Intelligent Data Graph makes that meaning usable in real time. Together, they ground AI in trusted, governed, continuously updated data, so agents can reason across entities, relationships, and interactions without drifting away from what's true.

The semantic graph explains what everything means. The Intelligent Data Graph makes that meaning operational in real time. Together, they ground AI in truth. They enable agents to reason across entities, relationships, and interactions while remaining anchored to high-quality, governed, real-time data.

What Did We Learn at Reltio as Customer Zero

As important as the technology is, the way we develop and test it is just as crucial. From the start, we believed we had to be Customer Zero for Reltio. That means we use our own platform to run our business before our customers rely on it for theirs.

Internally, we use Reltio to create an ontology that represents a 360 view of the Reltio business to see all our customers with their users, reconcile their complex usage with their entitlement information, and tie it to support and billing data. This information is integrated with our application and business process landscape across Sales, Support, Finance, Services, and Customer Experience. This allows us to experience the same integration and governance challenges that our customers face every day.

One of our early tests as Customer Zero came when we tried to reconcile usage data with entitlements and contracts across regions. The numbers didn't match. Different systems "owned" different answers. Our own teams were debating which report to trust.

Running that disconnected information through Reltio forced us to fix the data reconciliation issues, refine our data quality policies, and improve how we trace lineage across systems. Those changes didn't just fix our reporting. They became product capabilities that our customers now rely on to reduce weeks of manual reconciliation to real-time visibility.

Being Customer Zero is how we pressure-test our ideas about Context Intelligence. If our own agents and analytics can't answer hard questions about our business, we assume our customers' agents won't be able to either.

That experience has been invaluable. It has exposed edge cases, the messy realities of data, and the practical constraints that never appear in a product demonstration. It has also enabled us to innovate in collaboration with our customers.

As we have modernized our operations around an intelligent data layer and agents, we've followed the same path as global pharma companies building patient profiles; retailers unifying loyalty and commerce data; insurers connecting policy, claims, and risk data; and financial institutions harmonizing client, account, and risk data. We have not been spectators; we have been partners.

Looking ahead, we—at Reltio—are more convinced than ever that this combination of Context Intelligence and real-time operational data is what will separate the leaders from the laggards. Models will continue to improve, and new AI techniques will emerge. But the organizations that win will be the ones that invest in a shared context layer that makes every copilot, bot, and agent smarter, safer, and more aligned with how their business works.

The planned SAP acquisition reinforces our belief that this context layer will become foundational to how global enterprises operate, decide, and deploy AI at scale.

Reltio exists to build that context layer. We started by modernizing Master Data Management for the cloud. We extended it into a unified, graph-based view of entities, relationships, and interactions across the enterprise. Today, we are focused on making that context directly consumable by agents and intelligent applications, so that every decision, whether by a human or machine, can be made on trusted, connected, context-rich data.

Data Remains the Constant

Founding Reltio was, for me, an attempt to solve an old problem in a new way. Extending the boundaries with Reltio has been about something larger: helping enterprises reimagine what is possible when they take their data seriously and give intelligence the context it deserves.

My hope is that, through this book, you will see not only where the industry is headed but also how a deliberate investment in Context Intelligence

can advantage your organization for the future.

OceanofPDF.com

ABOUT THE AUTHORS



Manish Sood is the CEO, Founder, and Chairman of Reltio. Early in the shift to cloud and digital business, he saw a trend that many companies were missing: Data would need to power operations in real time, not just support analytics after the fact. That belief shaped Reltio from day one.

He founded Reltio in 2011 with a clear vision: connected, trusted, real-time data that moves as fast as the business. Under his leadership, Reltio has kept pushing that idea forward, unifying enterprise data, keeping it continuously current, and delivering it where teams need it most across customer engagement, supply chain, risk, compliance, and growth.

Today, many of the world's largest companies, including Fortune 500 leaders across industries, rely on Reltio to transform how they run with data. Manish's work is helping lay the data infrastructure for the agentic age, where AI agents and copilots will only be as effective as the trusted, real-time context

they can access. He is a frequent speaker, podcast guest, and writer on AI and enterprise data management.



Venkat Venkatraman is a leading global authority on digital business strategy. He is currently a Professor Emeritus at Boston University. He held the David J. McGrath Jr. Professorship of Management at Boston University's Questrom School of Business for over twenty-five years. Before this, he taught strategy at the MIT Sloan School of Management for nine years and at the London Business School for two years.

For the past forty years, Venkat's research and teaching have concentrated on how companies succeed in the Digital Age, where digital technologies influence and improve products, processes, and services. His highly regarded 2017 book, *The Digital Matrix: New Rules for Business Transformation through Technology*, was updated in 2023. His most recent coauthored book, *Fusion Strategy: How Real-Time Data and AI Will Power the Industrial Future*, was published by Harvard Business Review Press in March 2024. This book was one of the top three finalists for the Academy of Management's Best Book of the Year award and the Thinkers50 award.

Venkat has an undergraduate degree in Mechanical Engineering from the Indian Institute of Technology, Kharagpur, India; an MBA from the Indian Institute of Management, Calcutta, India; and a PhD in Strategic Management from the University of Pittsburgh. He is a distinguished alumnus of IIM Calcutta.

Venkat's doctoral thesis received the Academy of Management's AT Kearney Best Dissertation Prize, and the paper derived from the thesis, published in *Management Science*, is among the most cited strategy papers in the journal's history. According to Google Scholar, he ranks as one of the most highly cited researchers in management and is a leader in digital strategy, with over 65,000 citations. Recently, Stanford University recognized him as one of the top 2 percent of global scientists based on research impact. His work on business-IT alignment, published in *IBM Systems Journal*, was seen as a pivotal moment in our understanding of IT strategy and remains widely referenced.

Venkat writes papers for both academic publications and managerial audiences. His scholarly research has been published in prominent journals such as *Management Science*, *Strategic Management Journal*, *Information Systems Research*, *Academy of Management Journal*, and *Academy of Management Review*, among others. Over the last three decades, Venkat's articles for managers have appeared in the *Harvard Business Review*, *Sloan Management Review*, *California Management Review*, and *Financial Times*.

Venkat has consulted, delivered presentations, and conducted workshops globally for companies such as IBM, Ericsson, GE, BP, John Deere, Merck, GM, AWS, FedEx, Microsoft, McKinsey & Co., Philip Morris International, WPP, Sony, and Tesco. He served as a member of the Digital Technology Advisory Group at Canal+ in Paris for six years.

Venkat can be reached via LinkedIn and X @NVenkatraman.

ENDNOTES

- 1 “Data Quality: Best Practices for Accurate Insights,” Gartner, accessed January 20, 2026, <https://www.gartner.com/en/data-analytics/topics/data-quality>.
- 2 See <https://www.reltio.com/resources/case-studies/mercury-insurance> for more.
- 3 South Carolina Business Hall of Fame, “Edwin G. Seibels (1866–1954),” Knowitall.org (archived July 26, 2011), accessed January 20, 2026, <https://web.archive.org/web/20110726105725/http://www.knowitall.org/legacy/laureates/Edwin%20G.%20Seibels.html>.
- 4 For updated background information on MDM, see <https://www.ibm.com/think/topics/master-data-management> and <https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/master-data-management-the-key-to-getting-more-from-your-data>.
- 5 Those interested in knowing more about Aadhaar card, visit <https://uidai.gov.in/en/my-aadhaar/about-your-aadhaar.html>.
- 6 Vijay Govindarajan and Venkat Venkatraman, “The Next Great Digital Advantage,” *Harvard Business Review*, May–June 2022, <https://hbr.org/2022/05/the-next-great-digital-advantage>.
- 7 In late December 2025 and into early 2026, there was a lot of attention on ideas of “context graphs” popularized by Foundation Capital and primarily focused on how a new generation of start-ups could address this new market better than legacy software companies (<https://foundationcapital.com/context-graphs-ais-trillion-dollar-opportunity/>). Our thinking on agentic intelligence is consistent with those views, but we view it as broader than the software domain and believe that companies should develop their unique versions of “systems of context” for humans and AI to collaborate and execute.
- 8 Vivian Wu and Ramu Amanchi, “AbbVie Modernizes MDM Across Multiple Domains,” posted March 3, 2025, by Reltio, accessed January 20, 2026, YouTube, 22:02, <https://www.youtube.com/watch?v=UmgjBhG0B1Y>.
- 9 Simply Wall St., “Assessing AbbVie (ABBV) Valuation After Strong Multi-Year Shareholder Returns,” Yahoo Finance, January 5, 2026, accessed January 20, 2026, <https://finance.yahoo.com/news/assessing-abbvie-abbv-valuation-strong-170907693.html>.

- 10 See <https://db.financialexpressb2b.com/features/how-tata-steel-has-transformed-plant-and-mines-with-digital-twins> for an overview.
- 11 Abhishek Raval, “How Tata Steel Has Transformed Plant and Mines with Digital Twins,” Financial Express CIO, updated December 11, 2025, accessed January 20, 2026, <https://db.financialexpressb2b.com/features/how-tata-steel-has-transformed-plant-and-mines-with-digital-twins>.
- 12 “Walmart Data Ventures Shapes the Next Era of Insight-Driven Retail Innovation,” Walmart, October 29, 2025, accessed January 20, 2026, <https://corporate.walmart.com/news/2025/10/29/walmart-data-ventures-shapes-the-next-era-of-insight-driven-retail-innovation>.
- 13 “Technological Innovation,” Ant Group, accessed January 20, 2026, <https://www.antgroup.com/en/technology/>.
- 14 “We Have Built a Digital Society and We Can Show You How,” e-Estonia, accessed January 20, 2026, <https://e-estonia.com/>.
- 15 “The Origins of THINK,” IBM Heritage, accessed January 20, 2026, <https://www.ibm.com/history/think>.
- 16 “F1: The Movie (2025),” The Numbers, accessed January 20, 2026, [https://www.the-numbers.com/movie/F1-The-Movie-\(2025\)#tab=summary](https://www.the-numbers.com/movie/F1-The-Movie-(2025)#tab=summary).
- 17 Joel Shapiro, “Data Driven at 200 MPH: How Analytics Transforms Formula One Racing,” *Forbes*, January 26, 2023, accessed January 20, 2026, <https://www.forbes.com/sites/joelshapiro/2023/01/26/data-driven-at-200-mph-how-analytics-transforms-formula-one-racing/>.
- 18 “Amazon Study: Every 100ms in Added Page Load Time Cost 1% in Revenue,” Conductor, accessed January 20, 2026, <https://www.conductor.com/academy/page-speed-resources/faqamazon-page-speed-study/>.
- 19 “A 1-Second Delay Can Kill 7% of Conversions – Here’s the Data (2026 Guide),” WIRO Agency, January 15, 2026, accessed January 20, 2026, <https://www.wiro.agency/blog/how-a-1-second-delay-costs-you-a-7-drop-in-conversions/>.
- 20 Sarah Scher and Raymond Bunn, “Network Latency Costs More Than You Think, and We Have the Data to Prove It,” ArcGIS Blog (Esri), September 18, 2025, accessed January 20, 2026,

<https://www.esri.com/arcgis-blog/products/arcgis/administration/network-latency-costs-more-than-you-think-and-we-have-the-data-to-prove-it>.

- 21 Xinli Shang et al., “From Batch to Streaming: Accelerating Data Freshness in Uber’s Data Lake,” Uber Blog, December 11, 2025, accessed January 20, 2026, <https://www.uber.com/blog/from-batch-to-streaming-accelerating-data-freshness-in-ubers-data-lake/>.
- 22 Rafael Umann, “Real Time or Nothing: Why Low Latency Is the New Competitive Business Imperative,” *Forbes* (Forbes Tech Council), July 30, 2025, accessed January 20, 2026, <https://www.forbes.com/councils/forbestechcouncil/2025/07/30/real-time-or-nothing-why-low-latency-is-the-new-competitive-business-imperative/>.
- 23 For background discussions, please see <https://jakobnielsenphd.substack.com/p/think-time-ux>.
- 24 For an overview of Zara’s approach to fast fashion, see <https://medium.com/thedeephubzaras-fashion-revolution-through-data-insights-28f8fa728a0f>.
- 25 Kasra Ferdows et al., “Rapid-Fire Fulfillment,” *Harvard Business Review* 83, no. 11 (2005): 104–10.
- 26 An overview of the Join Life campaign can be found at <https://www.zara.com/us/en/z-join-life-mkt1399.html>.
- 27 Louis Thompsett, “Mastercard Secures 2nd Position in Top 100 Companies,” *FinTech Magazine*, October 13, 2025, accessed January 20, 2026, <https://fintechmagazine.com/news/mastercard-secures-2nd-position-in-top-100-companies>.
- 28 For an update on how Mastercard views payment trends with digital technologies, see <https://www.mastercard.com/global/en/news-and-trends/stories/2025/2026-payment-trends.html>.
- 29 For an update on Blackstone and data science, visit <https://www.blackstone.com/data-science/>.
- 30 Mark Harris, “Tesla’s Autopilot Depends on a Deluge of Data,” *IEEE Spectrum*, August 4, 2022, accessed January 20, 2026, <https://spectrum.ieee.org/tesla-autopilot-data-deluge>.
- 31 “Full Self-Driving (Supervised) Vehicle Safety Report,” Tesla, accessed January 20, 2026, <https://www.tesla.com/fsd/safety>.
- 32 Jimna Jayan, “Tesla’s Use of Data to Innovate in the Auto Industry: What Manufacturers Can Learn,” PromptCloud, October 1, 2024, accessed January 20, 2026, <https://www.promptcloud.com/blog/tesla-approach-to-automotive-data-solutions/>.

- 33 Mayo Clinic Press Editors, “AI, Big Data and Future Healthcare,” *Mayo Clinic Press*, July 1, 2025, accessed January 20, 2026, <https://mcpres.mayoclinic.org/research-innovation/ai-big-data-and-future-healthcare/>.
- 34 “Mayo Clinic Platform,” Mayo Clinic Platform, accessed January 20, 2026, <https://www.mayoclinicplatform.org/>.
- 35 Venkat Venkatraman, *The Digital Matrix: New Rules for Business Transformation Through Technology*, revised ed. (Wonderwell, 2023).
- 36 These machine agents would increasingly be a combination of virtual on the cloud and physical (such as robots and humanoids).
- 37 “DOT Penalizes Southwest Airlines \$140 Million for 2022 Holiday Meltdown,” US Department of Transportation, December 18, 2023, accessed January 20, 2026, <https://www.transportation.gov/briefing-room/dot-penalizes-southwest-airlines-140-million-2022-holiday-meltdown>.
- 38 *2025 Cultivating a Brighter Future Report* (Treasury Wine Estates Limited, 2025), accessed January 20, 2026, <https://www.listcorp.com/asx/twe/treasury-wine-estates-limited/news/2025-cultivating-a-brighter-future-report-3254067.html>.
- 39 “Cat® Production Measurement (CPM) Calibration Process for Excavators: Initial Calibration,” posted January 9, 2017, by Cat® Products, accessed January 20, 2026, YouTube, 11:05, <https://www.youtube.com/watch?v=M3gVHmhb4b0>.
- 40 Jordan Hart, “12 Pieces of Advice from Warren Buffett, from Parenting to Investments,” Yahoo Finance, May 6, 2025, accessed January 20, 2026, <https://finance.yahoo.com/news/12-pieces-advice-warren-buffett-163001603.html>.
- 41 Nino Paoli, “Walmart CEO Wants ‘Everybody to Make It to the Other Side’ and the Retail Giant Will Keep Headcount Flat for Now Even as AI Changes Every Job,” *Fortune*, September 27, 2025, accessed January 20, 2026, <https://fortune.com/2025/09/27/ai-ceos-job-market-transformation-walmart-accenture-sales-force/>.
- 42 “Already a Leading Car Maker, Mercedes-Benz Wants to Be an Innovative Software Company,” *The Software Report*, May 2, 2023, accessed January 20, 2026, <https://www.thesoftwarereport.com/already-a-leading-car-maker-mercedes-benz-wants-to-be-an-innovative-software-company/>.

- 43 Satya Nadella, “My Annual Letter: Thinking in Decades, Executing in Quarters,” LinkedIn, January 25, 2024, <https://www.linkedin.com/pulse/my-annual-letter-thinking-decades-executing-quarters-satya-nadella-7orpc/>.
- 44 Georgia Collins, “How Is Microsoft’s AI Strategy Driving Growth?,” *AI Magazine*, October 30, 2025, accessed January 20, 2026, <https://www.aimagazine.com/news/microsofts-new-growth-era-inside-satya-nadellas-ai-vision>.
- 45 Jessica Coacci, “Walmart CEO Says He Can’t Think of a Single Job That Won’t Be Changed by AI —Here’s How Today’s Workforce Can Prepare,” *Fortune*, September 30, 2025, accessed January 20, 2026, <https://fortune.com/2025/09/30/billion-dollar-retail-giant-walmart-ceo-doug-mcmillon-cant-think-of-a-single-job-that-wont-be-changed-by-ai-artifical-intelligence-how-employees-can-prepare/>.
- 46 Larry Dignan, “JPMorgan Chase’s Dimon on AI, Data, Cybersecurity and Managing Tech Shifts,” Constellation Research (Insights), June 11, 2025, accessed January 20, 2026, <https://www.constellationr.com/blog-news/insights/jpmorgan-chases-dimon-ai-data-cybersecurity-and-managing-tech-shifts>.
- 47 Dignan, “JPMorgan Chase’s Dimon.”
- 48 “Transforming Heavy Assets: Shell’s Digital Journey with Digital Twin and AI Technology,” Kongsberg Digital, December 8, 2024, accessed January 20, 2026, <https://kongsbergdigital.com/blog/transforming-heavy-assets-shells-digital-journey-with-digital-twin-and-ai-technology>.
- 49 James Elliott, “The IntelligentEngine,” Aerospace Manufacturing and Design, October 8, 2021, accessed January 20, 2026, <https://www.aerospacemanufacturinganddesign.com/article/the-intelligentengine>.
- 50 Janakiraman S. R., “How Does Technology Play a Role in Providing Passengers with a Positive and Enjoyable Flying Experience?,” Rolls-Royce, accessed January 20, 2026, <https://www.rolls-royce.com/country-sites/india/discover/2018/redefining-the-passenger-experience.aspx>.
- 51 Libby Hargreaves, “Ocado: Streamlining Supply Chains with AI Robotics,” Supply Chain Digital, April 15, 2025, accessed January 20, 2026, <https://www.supplychaindigital.com/technology/how-ai-is-redefining-ocados-robotic-fulfilment-system>.
- 52 “Ocado’s AI-Powered Robotic Arms: Levelling Up Efficiency in Online Grocery and Logistics,” Ocado Group, April 10, 2025, accessed January 20, 2026, <https://www.ocadogroup.com/newsroom/stories/ocado-robotic-arms>.

- 53 Reltio, “How Schneider Electric Is Driving Millions in Opportunities with Reltio,” Reltio Case Study, February 24, 2025, accessed January 20, 2026, <https://www.reltio.com/resources/case-studies/schneider-electric-case-study/>.
- 54 Reltio, “Republic Services Unlocks \$100M in Revenue by Unifying Data,” Reltio Case Study, accessed January 20, 2026, <https://www.reltio.com/resources/case-studies/republic-services/>.
- 55 Carnegie Mellon University, “CMU Partners with Moderna to Launch AI Academy,” *CMU News*, December 13, 2021, <https://www.cmu.edu/news/stories/archives/2021/december/moderna-ai-academy.html>.
- 56 Walmart Inc., “Walmart Offers a Glimpse into the Future of Retail at Consumer Electronics Show,” Walmart Corporate Newsroom, January 9, 2024, <https://corporate.walmart.com/news/2024/01/09/walmart-offers-a-glimpse-into-the-future-of-retail-at-consumer-electronics-show>.
- 57 “Leaders Underestimate Employees’ AI Use,” McKinsey & Company, March 4, 2025, accessed January 20, 2026, <https://www.mckinsey.com/featured-insights/week-in-charts/leaders-underestimate-employees-ai-use>.
- 58 Deborah Lovich et al., “GenAI Adoption Is Hard. Radical Employee Centricity Can Help,” Boston Consulting Group, August 5, 2025, accessed January 20, 2026, <https://www.bcg.com/publications/2025/genai-employee-experience-transformation>.
- 59 Stephen Dowling, “A ‘Toaster With a Lens’: The Story Behind the First Handheld Digital Camera,” BBC Future, December 8, 2025, accessed January 20, 2026, <https://www.bbc.com/future/article/20251205-how-the-handheld-digital-camera-was-born>.
- 60 “The Collapse of an American Retail Giant,” *The Economist*, October 20, 2018, accessed January 20, 2026, <https://www.economist.com/business/2018/10/20/the-collapse-of-an-american-retail-giant>.
- 61 Jeffrey Dastin, “Amazon Scraps Secret AI Recruiting Tool That Showed Bias Against Women,” Reuters, October 10, 2018, accessed January 20, 2026, <https://www.reuters.com/article/world/insight-amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK0AG>.
- 62 David F. Larcker and Brian Tayan, “Boeing 737 MAX,” The Harvard Law School Forum on Corporate Governance, June 6, 2024, accessed January 20, 2026, <https://www.corpgov.law.harvard.edu/2024/06/06/boeing-737-max/>.

- 63 Sean Michael Kerner, “CrowdStrike Outage Explained: What Caused It and What’s Next,” TechTarget, October 29, 2024, accessed January 20, 2026, <https://www.techtarget.com/whatis/featureExplaining-the-largest-IT-outage-in-history-and-whats-next>.
- 64 Patrick J. Kiger, “Flip Flop: Why Zillow’s Algorithmic Home Buying Venture Imploded,” Stanford Graduate School of Business, December 9, 2021, accessed January 20, 2026, <https://www.gsb.stanford.edu/insightsflip-flop-why-zillows-algorithmic-home-buying-venture-imploded>.
- 65 Tom Nawrocki, “Is UPI’s Rapid Growth Squeezing India’s Payments Market?,” PaymentsJournal, December 23, 2025, accessed January 20, 2026, <https://www.paymentsjournal.com/is-upis-rapid-growth-squeezing-indias-payments-market>.
- 66 Venkat Venkatraman, “IT-Enabled Business Transformation: From Automation to Business Scope Redefinition,” *MIT Sloan Management Review* 35, January 15, 1994, accessed January 20, 2026, <https://www.sloanreview.mit.edu/article/itenabled-business-transformation-from-automation-to-business-scope-redefinition/>.
- 67 Venkatraman, *The Digital Matrix*.
- 68 Vijay Govindarajan and Venkat Venkatraman, *Fusion Strategy: How Real-Time Data and AI Will Power the Industrial Future* (HBR Press, 2024).