

THE AGENTIC AI GENOME



Shannon Bell
Tom Jenkins
Steve Wagstaff

The Agentic AI Genome

By Shannon Bell, Tom Jenkins, and Steve Wagstaff

First publication, April 2026

Published by

Open Text Corporation

275 Frank Tompa Drive

Waterloo, Ontario, Canada

N2L 0A1

(519) 888-7111

info@opentext.com | www.opentext.com

How to Use This Book

Explaining the Trilogy

To provide readers with a comprehensive understanding of the evolving AI landscape, we have created a three-book series that can be read either as a complete trilogy or selectively, depending on your interests and familiarity with artificial intelligence.

Taken together, the trilogy follows the natural progression organizations experience as they move from experimenting with AI to operating intelligent systems at scale.

The first book in the series, *Enterprise Artificial Intelligence: Building Trusted AI in the Sovereign Cloud*, focuses on the foundation of trustworthy AI. It explains why the next era of AI innovation will belong to organizations that treat information as their most strategic asset. Drawing on decades of experience in secure, governed information management, the book explores how enterprises and governments can build AI systems that are not only powerful, but also secure, governed, and compliant.

The Agentic AI Genome, the second book in the series, builds on that foundation by introducing a new operating model for the enterprise. It presents a practical blueprint for designing and governing agentic AI at scale, demonstrating how agents, orchestration layers, enterprise information, and human oversight combine to form a governable architecture for intelligent workflows across the enterprise.

The third book in this series will take an even deeper dive into the implementation and orchestration methodology of agentic AI.

This Book: Designing the Agentic Enterprise

The purpose of *The Agentic AI Genome* is practical: to help organizations move from isolated AI pilots to a governed agentic operating model. Rather than treating AI as a collection of tools or experiments, this book introduces a structural model—the **Agentic AI Genome**—that explains how intelligent systems operate across an enterprise.

This is covered in five parts:

Part 1: Foundations – From Systems to Reasoning Enterprises (Chapters 1-4)

We begin by setting the stage for the agentic era, explaining how enterprise AI is evolving from tools that assist individuals to systems that can plan, act, and adapt across workflows. We show why trusted, governed information becomes the fuel of agentic execution—where content, records, operational data, and business network transactions form enterprise memory and evidence. From this foundation emerges the structural blueprint for the agentic enterprise: the **Agentic Genome**.

Part 2: Enterprise Functions and Agents Model (Chapters 5-7)

In this part of the book, we demonstrate how agentic AI becomes an operating model for intelligent work across enterprise functions and industry contexts. The unifying construct is the **Agentic Genome Map**: a consistent model for how agents, orchestration, enterprise information, and human command work together to deliver autonomous capabilities at scale. Rather than introducing different architectures for each function or sector, these chapters illustrate how the same genome expresses itself in different contexts. The structure remains consistent; what changes are the workflows, controls, data sources, and risk profiles specific to each domain.

Part 3: Technology and Methods (Chapters 8-10)

With the Agentic Genome defined, we turn to how it is built. This section translates the model into practical architecture, deployment patterns, and lifecycle discipline—showing how enterprise memory, agents, orchestration, and governance become a production system rather than an experiment. We examine the reference architecture for agentic enterprises, how these systems operate across hybrid and sovereign environments, and how the Agentic Development Lifecycle (ADLC) governs the creation, monitoring, and evolution of agents over time.

Part 4: Execution and Adoption (Chapters 11-13)

In this part, we focus on operationalizing the Agentic Genome inside real organizations. These chapters move from diagnosis to execution—first examining why most AI initiatives stall in pilots, then presenting a phased roadmap for scaling agentic AI into core workflows, and concluding with how leaders measure value, manage risk, and establish defensibility as autonomy increases. Together, these chapters shift the conversation from what agentic AI is to how organizations operationalize it at scale—with discipline, accountability, and measurable outcomes.

Part 5: What Comes Next? (Chapter 14)

This final section looks beyond architecture and adoption to the deeper transformation agentic AI creates inside organizations. As intelligent agents take on execution across workflows, the central question becomes how humans lead, govern, and collaborate with this new digital workforce. The future of the agentic enterprise will not be defined by automation alone, but by how successfully organizations balance autonomy with accountability, speed with trust, and technology with human judgment.

Where to Start Based on Your Role

Not every reader will approach this book with the same objectives. Some of you are responsible for enterprise architecture, while others are focused on financial oversight, workforce transformation, or public service delivery. The table below highlights recommended starting points depending on your role or area of responsibility.

If you are a...	Start with	Why it matters
CEO / Executive Leader	Part 1: Foundations; Chapter 14: Future of Human and AI Work	Understand how agentic AI changes operating models, leadership responsibilities, and enterprise competitiveness.
CIO / CTO / Chief Architect	Part 3: Technology and Methods	Learn how to design the technical foundation for agentic systems, including enterprise information substrates, orchestration layers, and governance architecture.
Chief Data Officer (CDO)	Part 1: Foundations; Part 3: Technology and Methods	Explore the role of enterprise information management, data governance, and trusted data pipelines in enabling safe agentic execution.
CFO / Finance Leader	Chapter 6, Finance and Audit; Chapter 13: Measuring Value and Defensibility	Understand how agentic AI transforms financial assurance, working capital management, and risk detection.
COO / Operations Leader	Chapter 6; Chapter 7: Operations and Supply Chain Agents; Part 4: Execution and Adoption	See how agentic systems optimize logistics, predictive maintenance, and cross-functional operational workflows.
CHRO / Workforce Leader	Chapter 6, HR and Talent Agents; Part 5: Chapter 14: Future of Human and AI Work	Learn how human–AI collaboration reshapes workforce skills, career paths, and organizational culture.
Chief Risk, Compliance, or Legal Officer	Chapter 6: Legal, Risk, and Records Agents; Chapter 7: Governance discussions throughout the book	Understand how auditability, policy enforcement, and explainability must be embedded into agentic systems.

If you are a...	Start with	Why it matters
Public Sector Leader / Agency Executive	Chapter 7: Federal, State/ Provincial, Municipal Agentic Workflows; Government case studies throughout the book (Part 2)	See how agentic architectures can improve citizen services, regulatory oversight, and administrative efficiency.
Enterprise AI Program Leader	Part 4: Execution and Adoption	Focus on the roadmap for scaling agentic AI from pilots to enterprise platforms.
Industry Specialists	Part 2: Chapter 7: Applying Agentic Capabilities to Private and Public Sectors	Discover how agentic architectures can safely scale operational efficiency and output while maintaining policy and regulatory compliance.

Using the Book as a Strategic Map

You may choose to follow the book sequentially, but many will find it most useful as a reference guide for designing your organization’s agentic strategy.

Because the Genome model applies across industries and sectors, you may also find value in exploring examples beyond your own domain. Many innovations emerge when organizations recognize that patterns developed in one sector—such as predictive maintenance in manufacturing or workflow automation in government—can be adapted to another.

While many AI books focus exclusively on commercial enterprises, this book deliberately includes examples from both private industry and government institutions.

The reason is simple: the architectural model for agentic AI is universal.

Applying the Genome Across Enterprise Domains

The book demonstrates how the Agentic Genome can be applied across the following enterprise functions:

Enterprise Domain	Agentic Capabilities
Executive Leadership	Strategic intelligence agents, board reporting automation, scenario simulation
Finance and Audit	Invoice automation, fraud detection, treasury forecasting
Human Resources	Employee onboarding orchestration, HR service agents, skills analysis
Sales and Marketing	Campaign optimization, dynamic pricing, lead qualification
IT Operations and Security	Autonomous incident response, vulnerability management
Operations and Supply Chain	Predictive maintenance, inventory balancing, logistics optimization
Legal, Risk, and Records	Contract analysis, regulatory monitoring, litigation discovery

Each domain uses the same architectural structure but applies it to different workflows and policy requirements.

Applying the Genome Across Public Sector Institutions

The same Genome map applies across public-sector organizations. Instead of departments driven by commercial objectives, government institutions often organize around mission-oriented services.

Public Sector Domain	Agentic Capabilities
Citizen Services	Case management automation, digital service orchestration
Justice and Courts	Document analysis, case workflow coordination
Public Administration	Identity governance, records management
Regulatory Agencies	Compliance monitoring, policy enforcement
Infrastructure and Transportation	Predictive maintenance and asset monitoring
Healthcare Systems	Diagnostic data analysis and operational scheduling

Why This Map Matters

Many organizations today deploy AI in fragmented ways—isolated chatbots, analytics tools, or automation scripts that operate independently of one another.

The Agentic Genome Map in this book helps organizations move beyond this fragmentation by providing a coherent architecture for intelligent execution. It allows you to visualize:

- Where agents operate within the enterprise
- Which workflows can be orchestrated across departments
- How governance and oversight must be structured
- How enterprise data must be organized to support agentic systems

In other words, the Agentic Genome transforms AI adoption from a series of experiments into a designed, scalable operating model.

Your journey starts here as, together, we move from insight to action.

Foreword

Over the course of our careers, we have witnessed three significant shifts in how organizations operate. The first was the rise of enterprise and SaaS software applications, which fundamentally changed how organizations work. The second was the cloud, which democratized access to infrastructure to run these applications. Today, we stand at the threshold of the third and most profound shift—one that is not about systems of record or systems of engagement, but about systems of reasoning and action.

This shift is the emergence of *agentic AI*.

Agentic AI does not simply answer questions or generate content. It plans, reasons, coordinates, and acts across systems. It turns intelligence into execution, moving organizations from using AI as a tool to operating AI as part of the enterprise itself.

This book marks the beginning of that new operating model.

Across industries, agentic AI has already begun to transform the foundational layers of organizations. Entry-level roles in customer support, human resources (HR) help desk, development, and quality testing are some of the first to evolve. Initially, agents served as copilots, augmenting human capability. But increasingly, agentic AI is becoming an autopilot, executing workflows with speed, consistency, and near-instant decision cycles.

These early use cases tend to be structured, repeatable, and governed by rules. They are not where strategy is defined or innovation thrives; they are where the foundations of the agentic enterprise are being laid. Yet, as agents take on more responsibility, leaders face a deeper architectural question:

If the bottom of the ladder is now automated, what becomes of the ladder itself?

And who climbs it—humans, agents, or both?

This is not simply a question about technology. It is a question about the future architecture of organizations and trajectory of human growth. The answer will impact how organizations are designed, how decisions are made, and how value is created in a world where intelligence can operate at machine speed.

What is emerging is not just a new class of technology, but a new organizational capability.

When paired with orchestrators that manage workflows, agents can break down complex tasks into thousands of micro-decisions, executed continuously and consistently, at speeds no human can match.

The result is extraordinary: enterprises that learn faster, respond faster, and operate on time scales that reshape competitive advantage.

This shift brings more than just operational change. It brings significant structural change. Just as email eliminated traditional mail rooms, agentic AI will reshape the hierarchy itself. Roles built to overcome human limitations will naturally evolve as agents assume greater parts of the organizational workload.

But this does *not* diminish the role of people. It elevates it.

In the agentic enterprise, humans become designers of activities, stewards of privacy and trust, and architects of outcomes. Humans will step out of routine tasks and into roles where judgment, empathy, strategy, and imagination matter more than ever. The defining leadership challenge is not how much autonomy to give machines, but how to design the boundaries within which autonomy operates—where humans must remain in the loop, and how agents can expand the loop around them to extend human capacity.

To build this future responsibly, we must prioritize one fundamental requirement: trusted, complete, and governable data.

Agentic AI cannot reason or act with confidence unless it is grounded in the full spectrum of organizational knowledge—human-generated content, machine-generated signals, and transactional systems of record. Our enterprise applications, including Customer Relationship Management (CRM), Enterprise Resource Planning (ERP), Service Management, and Supply Chain Business Networks, have spent decades encoding the rules of our operations. In an agentic world, these systems become the orchestrators of tomorrow, coordinating agents the way they once coordinated people.

When we reimagine the enterprise through this lens, something remarkable happens.

A global corporation of 100,000 people may reveal only a few thousand distinct roles, each of which can be mapped to an agent. The millions of workflows that drive daily operations collapse into a few hundred orchestrators. Complexity becomes simple. Scale becomes manageable. And the organization becomes not only faster, but fundamentally smarter.

The Agentic AI Genome describes this transformation conceptually.

It explores how enterprises can move from isolated AI pilots to an agentic operating model. It examines how agents, orchestrators, and governed information together form a new enterprise “genome.” It shows where real productivity, resilience, and strategic advantage emerge when intelligence is embedded into the fabric of how organizations run. And it highlights the use cases where productivity breakthroughs will emerge.

We are entering an era in which bottlenecks created by human handoffs, application silos, and rigid hierarchies will no longer define the limits of performance. Instead, we will be limited only by the quality of our data, the clarity of our purpose, and the boldness of our vision.

Those who embrace agentic AI early and thoughtfully will define the next generation of performance, resilience, and value creation. The competitive frontier is no longer infrastructure alone—it is orchestration at machine speed.

The future enterprise will not simply use AI. It will be built with AI, built for AI, and ultimately built as AI.

The Agentic AI Genome is your guide to that future.

About the Authors



Shannon Bell

Shannon Bell is the Executive Vice President, Chief Digital Officer, and Chief Information Officer for OpenText, responsible for the company's IT and digital systems, data platforms, networks and communications, commercial and corporate cloud operations, as well as its security and compliance. She is an accomplished IT leader with over 25 years of international experience in technology transformation and large scale integrations. Prior to OpenText, she held a senior leadership role at Rogers Communications, spearheading the technology integration of the Shaw acquisition. Her career has included roles at Amdocs, NewStep Networks, MetaSolv Software, Axiom Systems, and Newbridge Networks. Shannon holds an MBA from the University of Surrey and an undergraduate degree from Carleton University. Shannon is co-author of the first book in this series: *Enterprise Artificial Intelligence: Building Trusted AI in the Sovereign Cloud*.



Tom Jenkins

One of Canada's leading experts on innovation and digital technologies, Tom Jenkins is the Chair of OpenText Corporation, the largest software and cloud company in Canadian history and one of the most successful internet companies in the world. Tom has served or continues to serve on the boards of OpenText Corporation, Manulife Financial, Thomson Reuters, TransAlta Corporation, BMC Corporation, and Slater Steel. He also served as the Chair of the National Research Council of Canada. He received his commission as an officer in the Canadian Armed Forces and as the Honorary Colonel of an infantry regiment and a fighter squadron in the Canadian Armed Forces. He was the 10th Chancellor of the University of Waterloo, and he was inducted as a Companion into the Canadian Business Hall of Fame. Tom is a recipient of the Federal Republic of German Order of Merit (Knight's Cross), and he is an Officer of the Order of Canada. Tom has authored 10 business books on digital innovation, and co-authored *Ingenious: How Canadian Innovators Made the World Smarter, Smaller, Kinder, Safer, Healthier, Wealthier, and Happier*, with David Johnston, former Governor General of Canada. Tom is also co-author of *Enterprise Artificial Intelligence: Building Trusted AI in the Sovereign Cloud*.



Steve Wagstaff

Steve Wagstaff is Senior Vice President of Marketing at OpenText, where he helps shape the company's global market narrative and strategic positioning. His work focuses on defining how enterprises understand and operationalize agentic AI, reframing it from a technology trend into a fundamental shift in how organizations design work, decision-making, and value creation. With more than 35 years of experience spanning technical sales, value engineering, and business transformation, Steve has built a career at the intersection of innovation and business strategy. At OpenText, he is recognized for establishing the company's Value Engineering discipline and for leading Worldwide Enterprise Presales for over a decade, guiding organizations in translating complex technologies into measurable strategic and financial outcomes. Earlier in his career, Steve was a pioneer in the boutique personal computer industry. He founded and served as President of Ares Microdevelopment, a Michigan-based high-performance PC manufacturer known for its "military-style" build quality and the industry's first Lifetime Warranty. Under his leadership, Ares became a frequent recipient of *PC Magazine's* coveted Editor's Choice Awards.

Acknowledgements

The authors would like to thank the following contributors:

Michael Acedo, Stephane Anastassiades, Mahitab Andelsalam, DeeDee Andrews, Ayman Antoun, Matt Arcaro, Shannon Berg, Michelle Berry, Todd Cione, Paul Duggan, Joe Dwyer, Magali Germain, Pascale Guivier, Adam Hennessy, Sandra Herber, Michelle Landaverde, Mark L'Heureux, Stephen Ludlow, Sarah Loat, Mihaela Lucaci, Muhi Majzoub, James McGourlay, Bitu Houshmand Rabiee, Steve Rai, Sunnie Rothenburger, Zamantha Sanchez, Scott Schultz, Tim Spadzinski, along with Elizabeth Chestney-Hanson, editor, and Kevin Sy for layout and design.

Contents

How to Use This Book	3
Foreword	9
About the Authors	11

Part 1: Foundations – From Systems to Reasoning Enterprises

Chapter One

The Rise of the Agentic Enterprise	16
---	----

Chapter Two

Information as Fuel: Content Systems Empower AI	31
--	----

Chapter Three

Sovereignty, Accountability, and Control	46
---	----

Chapter Four

From Org Charts to Orchestration	57
---	----

Part 2: Enterprise Functions and Agents Model

Chapter Five

What Is an Agentic Genome Map? (and Why You Need One)	76
--	----

Chapter Six

Applying Agentic Capabilities to the Enterprise	95
--	----

Executive and Board Agents	96
----------------------------	----

Finance and Audit Agents	98
--------------------------	----

HR and Talent Agents	102
----------------------	-----

Sales and Marketing Agents	105
----------------------------	-----

IT, Security, and Compliance Agents	108
-------------------------------------	-----

Operations, Supply Chain, and Facilities Agents	112
---	-----

Legal, Risk, and Records Agents	115
---------------------------------	-----

Chapter Seven

Applying Agentic Capabilities to the Private and Public Sectors	121
--	-----

Banking and Finance	124
---------------------	-----

Food and Beverage Manufacturing	128
---------------------------------	-----

Transportation and Logistics	133
------------------------------	-----

Automotive	137
------------	-----

Healthcare	141
------------	-----

Pharmaceutical	146
Process Manufacturing	151
Oil and Gas	156
Energy and Utilities	160
Public Sector – Federal	165
Public Sector – State/Provincial	170
Public Sector – Local/Municipal	174
Part 3: Technology and Methods	
Chapter Eight	
Reference Architecture for Agentic Enterprises	182
Chapter Nine	
Deployment Models: Multi-Cloud, Hybrid, & Sovereign	195
Chapter Ten	
The Agent Lifecycle	208
Part 4: Execution and Adoption	
Chapter Eleven	
From Pilots to Platforms	224
Chapter Twelve	
A Phased Adoption Roadmap	239
Chapter Thirteen	
Measuring Value and Defensibility	253
Part 5: What Comes Next?	
Chapter Fourteen	
The Future of Human and AI Work	266
Appendices	
Endnotes	276
Glossary	281
Works Cited	292
Index	298

Part 1

Foundations – From Systems to Reasoning Enterprises

01

Chapter One

The Rise of the Agentic Enterprise

This chapter is about the transformation of AI from a tool that supports decisions to a capability that executes work. We examine how enterprises have progressed from predictive and generative AI to agentic systems that can plan, act, and adapt across complex business processes. This shift signals the rise of a new operating model—one in which intelligence becomes embedded in the fabric of how the enterprise runs.

From Static AI to Agentic Intelligence: The Paradigm Shift

The first wave of enterprise AI (EAI), from the late twentieth to the early twenty-first centuries, focused on insight. The term AI encompassed everything from predictive and advanced analytics to intelligent search and even automation. Organizations used machine learning models to forecast demand, detect customer churn, or classify customer segments. These systems improved decision quality, but they did not change how work flowed through the enterprise.

The second wave, well into the second decade of the twenty-first century, brought generative AI (GenAI) and Retrieval-Augmented Generation (RAG)—a process that improves the accuracy and reliability of Large Language Models (LLMs) by anchoring them in specific and relevant data sources. Knowledge search, document summarization, and content creation improved dramatically. AI became a powerful interface for enterprise information—responsive, conversational, and increasingly capable.

In all these applications, AI provides single-output support: a prediction, a classification, or generated content, typically in response to a human prompt or query. Many of these systems rely on the enterprise having a trusted, well-managed data foundation, a critical requirement we stress in the foundational book in this series, *Enterprise Artificial Intelligence: Building Trusted AI in the Sovereign Cloud*.

While powerful, this second-wave model is inherently limited. Every task must be framed with the right prompt. AI is not planning, reasoning, or taking initiative. A human is driving the workflow: collecting data, deciding what to do next, executing, reviewing, and iterating. This “human-in-the-loop” paradigm, while essential for control and oversight, has quickly become a bottleneck for scale, agility, and real-time responsiveness.

A new paradigm is emerging: agentic AI.

Unlike earlier AI systems designed to generate insight, agentic systems participate directly in execution. Rather than simply responding with a prediction or piece of content, these systems operate with significant initiative and flexibility, whether acting autonomously or in collaboration with humans.

What Is Agentic AI?

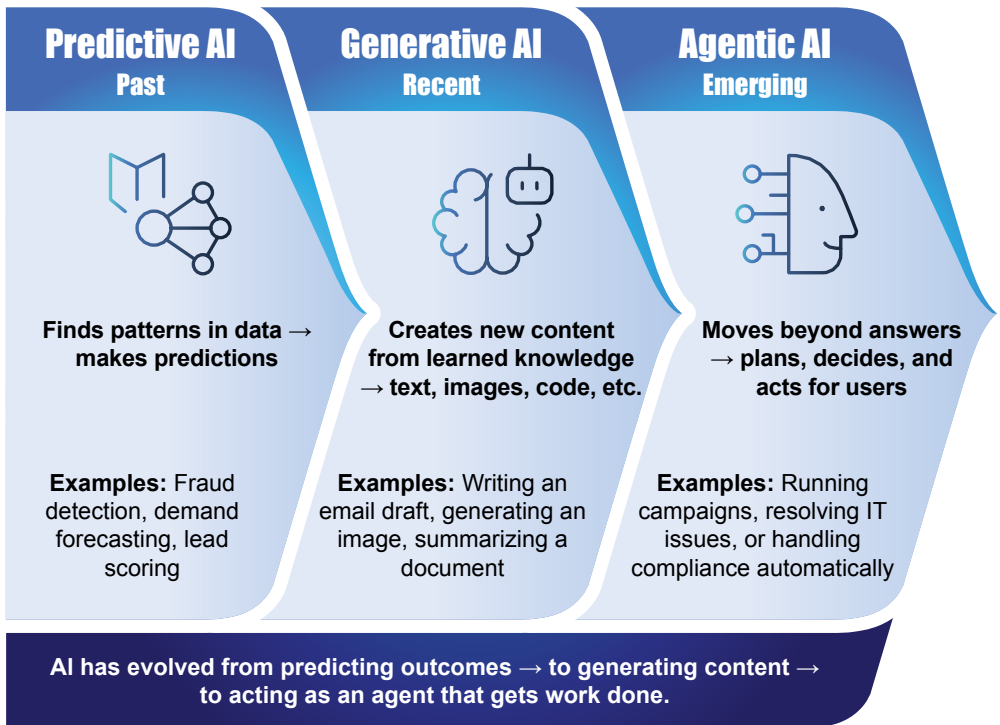
Agentic AI represents a fundamental departure from previous generations of artificial intelligence. While traditional AI systems require explicit programming for each task, and generative AI focuses on content synthesis and creation, agentic AI operates with autonomous, goal-seeking behavior. It adapts strategies and executes complex business processes independently.

Think of it as a highly capable digital employee. You can assign it a high-level objective—for example, “process this customer refund”—and it will autonomously drive the workflow: retrieving relevant data, deciding whether conditions are met, updating systems, notifying stakeholders, and escalating if needed. It figures out the steps, executes the plan, and adapts as needed to get the job done.

In technical terms, an agentic system consists of multiple AI agents, each responsible for a specific task, working through an orchestration layer that coordinates their actions.

This differs significantly from previous paradigms:

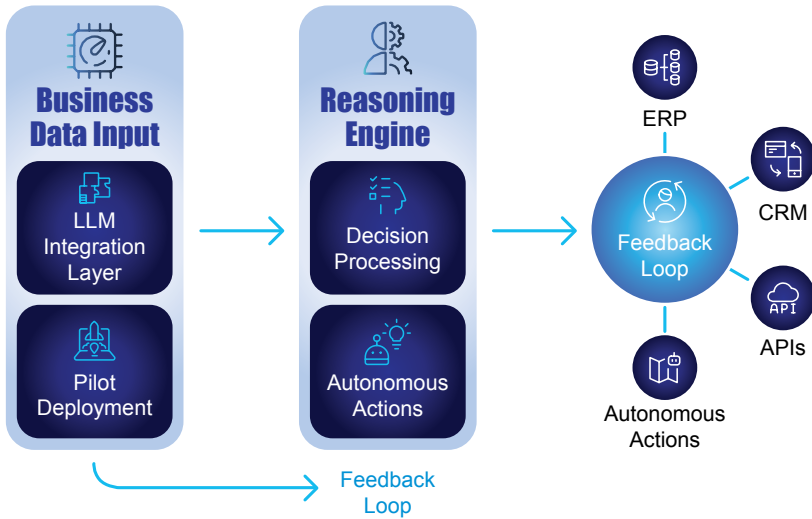
- Unlike traditional AI, which follows rigid, pre-designed rules and breaks when it encounters unexpected variables, agentic AI uses reasoning to adapt to new environments without constant reprogramming.
- Unlike generative AI, which typically provides a single output (text or image) in response to an immediate prompt, agentic AI maintains objectives across multiple interactions. It can independently plan and execute multi-step processes to achieve specified goals.



From Predictive to Agentic AI

The result isn't just a prediction or summary; it's a completed task or workflow that has undergone planning, execution, and coordination. This shift allows organizations to deploy end-to-end automated processes that combine data analysis with strategic execution without requiring human intervention at each decision point.

This autonomy is based on the agent's understanding of the environment, its assigned goals, and its learned experiences. It can analyze data, assess risks, and choose the best course of action to achieve its objectives.



How Agentic AI Works

Its value is iterative. Agentic AI is adaptable, learning from its interactions to improve performance over time. As the agents encounter new situations and receive feedback, they refine their decision-making processes and become more adept at achieving their goals.

This shift from static to dynamic AI reflects a deeper transformation: from AI as a tool for search and summarization to AI as a digital workforce, capable of deliberate, autonomous action without constant reprogramming or human oversight.

The journey from second-wave summarization to third-wave autonomous action is best observed in high-stakes legal and operational environments. In the following case study, Bosch illustrates this transition by moving beyond simple keyword search toward an intelligent discovery engine that plans and executes complex data investigations at scale.

Bosch



We can see a significant return on investment (ROI) from using AI capabilities in legal decision-making. It is too soon to put an exact number on it, but I estimate cost savings in the millions.

– Franziska Fuchs, Vice President of eDiscovery, Bosch

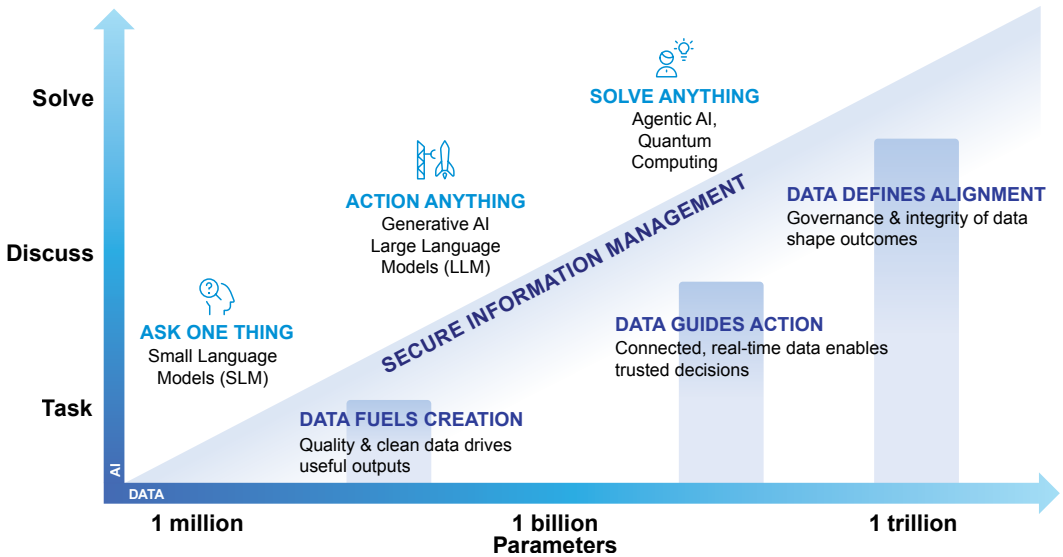
A global leader in technology and services, Bosch shapes universal trends in automation, electrification, digitalization, and connectivity. Its strong industrial presence drives innovation to improve processes, including optimizing legal operations. At the company, legal disputes distracted from its core mandate of driving continuous innovation.

In any legal dispute, understanding the data early and deeply is crucial but challenging. If the legal department could uncover relevant information and supporting evidence much earlier in the process, it would help ensure data-driven legal and business decisions. Technology could help Bosch make better and faster decisions by quickly processing and analyzing large amounts of internal data. The resulting knowledge could then shape how a legal case proceeds. Given the high stakes and long-term nature of legal cases, Bosch looked for a reliable technology partner, focused on consistent delivery and innovation.

Bosch is using GenAI to transform cumbersome legal tasks and validate its ability to accelerate investigations and early case assessment, reduce costs, and enhance decisions—to ensure high-quality outcomes. Using AI, Bosch has transformed legal workflows, enabling faster data-driven decisions and proactive investigations. Continuous data query capabilities ensure no critical information is missed. AI integration and training empower legal teams to deliver superior service. Legal teams learn to integrate AI into everyday tasks, shifting from manual review to strategic decision-making.

By delegating the labor-intensive scout work of data synthesis to an intelligent system, Bosch has effectively created a digital workforce that handles the first 90% of the investigative lifecycle. This allows their human experts to focus exclusively on high-level strategy and judgment, proving that the rise of the agentic enterprise is ultimately about reclaiming human capacity for the decisions that matter most.

The Rise of Agentic AI



The Evolution of Agentic AI

As AI and workforce paradigms evolve, the following forces are converging inside the enterprise:

- **The proliferation of data.** As enterprises automate more processes, including CRM, ERP, supply chain, customer interactions, documentation, and logs, the volume and variety of available data have increased exponentially. This data enables agents to reason and act across business domains—if that information is governed, trusted, and accessible.
- **The maturation of AI models and tooling.** LLMs, reasoning capabilities, and orchestration platforms have matured to the point where autonomous agents are technically feasible and reliable.
- **The demand for real-time, scalable decision-making.** Markets move quickly, and customer needs are evolving. Enterprises increasingly require systems that can react and adapt in real time, not just generate reports overnight. Execution speed is becoming a competitive differentiator.
- **The need to eliminate human-in-the-loop bottlenecks.** Human review and intervention give control, but also limit scale, speed, and consistency. Agentic systems can offload low-value, repetitive, or structured decision workflows, freeing humans to focus on strategic, creative, or higher-stakes tasks.
- **The imperative to bridge the “Pilot Gap.”** Organizations have invested heavily in AI experiments, but few have translated pilots into enterprise platforms. The gap between experimentation and institutional capability is now the central leadership challenge. First movers will create a new type of enterprise that will put slow movers out of business.

In short, this transformation marks a shift from AI as a tool to AI as an operating capability.

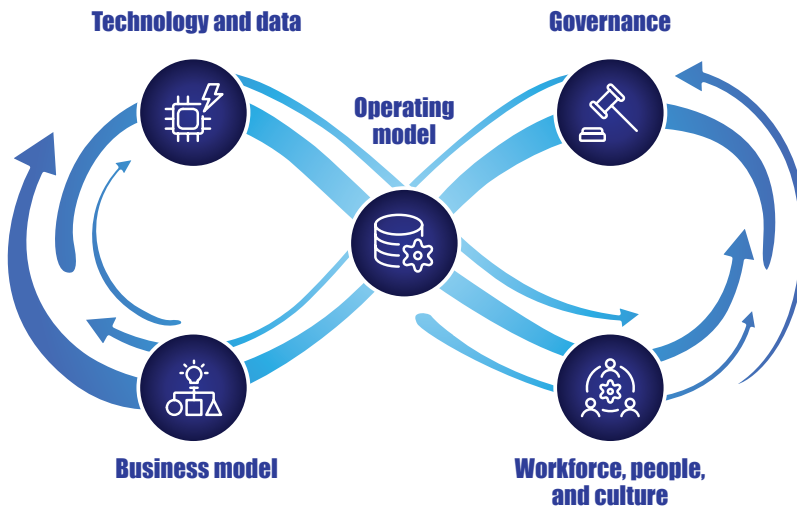
What distinguishes agentic AI is not only autonomy; it is orchestration. Multiple agents operate in concert, coordinating through workflow and policy layers that bind intelligence to execution. In this model, content platforms and enterprise systems become the memory and control plane for intelligent action. AI moves from the edge of the enterprise into the fabric of how work is done.

This is the rise of the agentic enterprise.

The Agentic Enterprise: Capabilities and Possibilities

We have just explored how agentic AI is transforming the way businesses approach automation, decision-making, and problem-solving. Humans define objectives, constraints, and accountability. Agents execute within those boundaries. By delegating complex, multi-step processes to AI agents, organizations and government agencies can better focus on strategic initiatives, innovation, and customer relationships while achieving unprecedented operational efficiency.

What does a truly agentic enterprise look like in practice? Imagine an organization where AI agents, not just human employees, carry out end-to-end workflows across systems, around the clock, working in concert.



The Agentic Enterprise Will Unlock Transformation Across Key Pillars¹

McKinsey & Company examines this shift in their analysis of the agentic organization:

As agents take on execution, people will increasingly define goals, make trade-offs, and steer outcomes. This will change how companies plan for a hybrid workforce, who they hire (or borrow), how they deploy human or AI talent, and how they measure success. HR systems not only track human employees but also are a repository of agents and agentic workflow.²

Below are examples of these capabilities in action:

Autonomous workflows: Agents complete multi-step decision processes. For example, an AI agent could handle the entire invoice-approval process: starting with retrieving invoice data, validating against purchase orders, checking budget constraints, routing to the correct approver, logging the decision, and triggering payment—all without human intervention, aside from the final approval.

Cross-system orchestration: Agents coordinate across CRM, ERP, finance, procurement, and support systems. For example, when a supply-chain issue arises, an agent can detect it, trigger alternative sourcing workflows, alert stakeholders, and perform financial impact analysis.

Continuous operations (24/7): Agents monitor system performance, detect anomalies, remediate common problems, and escalate only for exceptional cases, achieving round-the-clock resilience without added human headcount.

Dynamic scaling and adaptability: Agents adapt to market dynamics. As volume shifts, agentic systems can autoscale workflows, re-prioritize tasks, and reallocate resources.



These capabilities deliver significant benefits:

- **Agility and scalability:** Enterprises can respond fast to changing conditions with agents that scale elastically.
- **Faster execution:** Workflows and processes that once took days or hours can now be completed in minutes or seconds. In fact, industry experience suggests that agentic AI can accelerate business processes by 30-50%.³
- **Error reduction and higher-value human work:** Autonomous agents reduce manual repetitive work, lowering error rates and freeing human labor for higher-value tasks.
- **Strategic transformation:** The enterprise moves away from discrete departments and manual handoffs, becoming an integrated, AI-driven system. Workflows, data, decisions, and orchestration all operate as a cohesive digital unit.

The agentic enterprise achieves new levels of operational resilience, speed, cost-efficiency, and flexibility, opening the door to business models and ways of working that were previously unachievable.

The impact is not incremental. It is structural. Departments become capability domains. Hand-offs give way to orchestration. The enterprise begins to behave less like a set of applications and more like a coherent operating system.

This is not digital transformation. This is *institutional memory and reasoning at enterprise scale*.

While many organizations struggle to bridge the pilot gap, the following case study of a major European airport demonstrates how a high-volume, regulated environment can move beyond static IT support to an anticipatory operating system where agents orchestrate 24/7 resilience across millions of passenger touchpoints.



A European Airport



This European Airport has long been a hub of operational complexity and service expectations. Serving more than 28 million passengers annually, the airport faced the challenge of providing consistent, reliable service across a diverse set of users and systems. To modernize its IT service management, it adopted a cloud-hosted service management platform that unified incident, problem, change, and configuration workflows into a single, integrated system. With this platform, the airport significantly improved resolution times, increased visibility into IT operations, and empowered users with self-service capabilities.


This modernization serves as the jumping-off point for a more intelligent, autonomous operating model. Rather than simply automating ticket routing, an agentic orchestrator now continually ingests real-time telemetry from IT systems, infrastructure sensors, and passenger-facing applications. Diagnostic agents proactively detect patterns that signal emerging service degradation—such as network congestion at peak travel times—initiating workflows without human intervention. The system automatically schedules remediation tasks, updates configuration records, and triggers knowledge agents to update documentation with complete traceability.

Agentic AI transforms the airport's IT ecosystem from reactive service support into a coordinated, adaptive engine of operational resilience. Disruptions can be detected and addressed in shorter timeframes, infrastructure lifecycles are extended through autonomous condition-based maintenance, and human experts are freed to focus on strategic improvements rather than repetitive tasks. By anchoring agentic intelligence on a trusted content foundation and orchestrating across domains at machine speed, the airport exemplifies how complex, regulated enterprises can evolve into responsive, anticipatory systems that uphold experience, safety, and performance in an increasingly dynamic world.

Challenges, Risks, and the Need for Governance

The rise of the agentic enterprise is not inevitable. Constrained by architecture, governance, and leadership, it's not without its challenges.

A recent MIT report reveals the stark reality of the "GenAI Divide": "Despite \$30–40 billion in enterprise investment into GenAI... 95% of organizations are getting zero return. Just 5% of integrated AI pilots are extracting millions in value, while the vast majority remain stuck with no measurable P&L impact."⁴



The GenAI Divide—from Pilots to Production⁵

Against this backdrop of challenges in GenAI, agentic AI has a tough hill to climb. Significant obstacles and risks stand in the way, especially given the security, compliance, and data governance implications of EAI.

Structural challenges involve data silos, legacy systems, and dated infrastructure. Many enterprises remain burdened with legacy tools, disconnected systems, inconsistent data models, and poorly integrated workflows. Without a unified data foundation that ensures high-quality, consistent, governed, and trusted data, agentic systems may operate on incomplete or incorrect information.

The first book in this series emphasizes that trusted data is not optional. Even more, integrating agents deeply into enterprise systems requires careful architecture and strong data governance.

Beyond structural challenges, agentic autonomy can bring unpredictability. Agents making independent decisions can lead to unintended consequences, cascading errors, or actions taken out of context. As researchers warn, agentic systems introduce risks of tool misuse, cascading action chains, and emergent behaviors that are difficult to foresee.⁶

From a compliance, ethical, and regulatory standpoint, enterprises may struggle to ensure transparency, explainability, accountability, and auditability. In regulated domains like finance, healthcare, and supply chain, “letting the agent decide” is often unacceptable without strict guardrails.

Organizational readiness remains a major hurdle as well. Becoming an agentic enterprise is not just a technical shift or challenge; it is a cultural one. Human employees must evolve from *performing* tasks to *supervising* agents. This requires thoughtful change management, as fear of job loss or uncertainty about new roles can impact employees’ willingness to adopt agentic AI. Without upskilling and a clear plan for human–AI collaboration, adoption may not be successful.

Finally, governance and security present roadblocks. Existing protocols for trusted data and security are a necessary foundation, but they are no longer sufficient. Agentic AI demands new governance models that define the boundaries for autonomous behavior, specify when humans must be in the loop, ensure proper audit trails, decision-logging, and control access, and safeguard data privacy. Without robust governance, enterprise adoption risks failure, security breaches, or unacceptable mistakes.

The path to the agentic enterprise is challenging, but the rewards are enormous. The transformation is about more than plugging in more intelligent AI; it’s about rearchitecting data, systems, processes, culture, and governance.

Implementing agentic AI raises a leadership question as much as a technical one: *How do you scale intelligent execution without losing control?*

If the GenAI divide is primarily a gap in measurable P&L impact, then bridging it requires moving beyond experimental prompts toward the rigid structural discipline demonstrated by Volkswagen Finance China (below)—where governing an 800% explosion in data was treated not as a storage problem, but as a mandatory architectural prerequisite for autonomous compliance.

VW Finance China



Volkswagen Finance China

Volkswagen Finance (China) Co., Ltd. is China's first wholly foreign-owned auto finance company. Volkswagen Finance China has been committed to providing Chinese customers with advanced auto finance products and high-quality services.

As a result of the rise of affluence in the Chinese market and the appeal of foreign car brands, there has been an increase in the purchase of cars and a corresponding demand for automotive financial services. Business for Volkswagen Finance China has been growing rapidly and the corresponding data has increased by 800%. In the midst of this deluge of data, the company is required to comply with regulations in both its host country, China, and its country of origin, Germany. These requirements include the need to store related documents for seven years in both paper and digital formats in support of the electronic disbursement process.

The company turned to an expansive information management solution to give it the agility it requires to manage compliance data and processes—both now and into the future. The solution helps to manage compliance using integrated records management, archiving, search and e-Discovery. The solution can handle large amounts of data and integrates with other applications like email and its retail finance suite to minimize the risk and cost associated with compliance, while maximizing insight and efficiency across consolidated data sources. As well as transparency into performance, the organization has been able to find new value in this solution by reducing the time required to find documents by 50 percent, significantly improving productivity and time-to-compliance.

VW Finance China’s journey underscores a vital lesson for the agentic era: autonomy is only as safe as the records management system that constrains it. By cutting document discovery time in half and automating the preservation of high-stakes financial data, the company has built the exact sovereign data foundation required to scale agentic workflows without risking regulatory breach.

From Pilots to the Agentic Operating Model

When we talk about the “rise” of the agentic enterprise, what do we really mean? It is not a single moment in time, but a transformational journey.

Many organizations are still at the pilot stage with proof-of-concept agents solving narrow tasks. These initiatives demonstrate possibility, but they often amount to little more than advanced automation; they do not change how the enterprise runs. Others may have deployed small-scale agentic automation implementations, but they are isolated, disconnected, or not deeply integrated into enterprise workflows.

The target state is a full-scale *agentic operating model*—an enterprise where autonomous agents, orchestration layers, and human oversight combine to form a cohesive, scalable, resilient, and flexible operating system. In this model, the data must be trusted, governed, and integrated with agents embedded across business functions. Workflows should be dynamic and adaptive, and humans must focus on oversight, governance, and exception handling.

For executives, this shift matters because, when done well, it promises competitive advantage, operational resilience, and the ability to handle complexity and growth at scale. It represents the future of human and digital resource management.

This book is about making the shift deliberately. It will show you how to build the data, content, and organizational foundation for a successful agentic AI transformation. It builds on the fundamentals of trusted data and secure AI, extending them into the agentic era.

Specifically, it introduces the concept of the “Agentic Genome Map”—a framework to map how agents, orchestrators, content, and governance together define the operating model of the modern enterprise.

Executive Implications

CAIO

The shift from advisory AI to execution creates a new leadership imperative. The CAIO defines the governance architecture—risk classification, oversight, and coordination—that enables agents to operate safely at scale and prevents uncontrolled pilot sprawl.

CIO

The shift from AI that advises to AI that executes requires treating agentic systems as an operating model—architected for orchestration, trusted enterprise memory, and controlled integration across core platforms.

CFO

Agentic AI changes the economic equation from incremental productivity to structural operating leverage, but only if organizations close the pilot gap and tie autonomous workflows to measurable P&L outcomes.

CHRO

As work moves from humans performing tasks to humans supervising digital labor, HR must redesign roles, skills, and change management for a hybrid workforce where agents execute and people govern.

CDO

Execution-grade AI amplifies the consequences of poor data quality and weak governance; agentic scale depends on trusted, accessible, and auditable information foundations.

COO

Operational advantage comes from replacing manual handoffs with orchestrated end-to-end execution, enabling faster cycle times, 24/7 resilience, and adaptive responses without adding headcount.

Together, we will explore how to design, govern, and deploy agentic systems at scale, providing use cases across workflows, job functions, and industries to help you create an agentic organization.

The goal is not to run more pilots.

The goal is to build the enterprise that can operate with intelligence at scale.

Let's begin.

Chapter Two

Information as Fuel – Content Systems Empower AI

In an agentic enterprise, intelligence doesn't live in AI models alone. It lives in the *information* those models can read, reason over, and act upon.

This chapter explains why content systems form the memory, trust layer, and control plane for agentic AI. We move beyond viewing data as merely input, shifting toward information as the operating fuel of intelligent action. This transition requires content architectures to evolve to support orchestration, governance, and human-in-command design. Only then can you design an agent-ready information foundation for your enterprise that enables scale without sacrificing trust, sovereignty, or control.

The Nature of Data and Content in the Agentic Era

The first book in this series, *Enterprise Artificial Intelligence: Building Trusted AI in the Sovereign Cloud*, establishes “trusted data” as the foundation of enterprise-grade AI: data quality, governance, security, and compliance are non-negotiable.

In an agentic AI world, this trusted data paradigm must be extended. Data becomes the operational foundation upon which intelligent agents operate, reason, and act.

As a result:

- Enterprises must leverage both structured and unstructured information, enriched with metadata, lineage, context, and access policies.
- Content is no longer “reference material.” It is the memory, context, evidence, and audit log of agentic workflows.
- Data integrity—accuracy, currency, completeness—is operationally critical, because agents act on what they observe. Data errors become business errors.
- Governance and security are no longer just about protecting data; they are about governing intelligent behavior

In the agentic enterprise, information becomes a living, governed medium of execution, rather than a passive repository.

Before an agent can reason across a lifecycle, the enterprise must first solve the fragmented memory problem. In the following case study, Buncombe County illustrates this shift by transforming static, duplicated paper records into a single, governed data storehouse that serves as the essential cognitive substrate for automated administrative action.



Buncombe County



Streamlining Processes at Buncombe County

Buncombe County, a local level of government in North Carolina, needed to enhance information sharing and reduce the amount of paper created and stored. With employees in almost two dozen different areas, Buncombe County wanted a way to streamline administrative processes and improve customer service.

In Human Resources, for example, all processes were paper based and each department had their own copies of their employee files; payroll had copies and personnel had the official employee record. There was a tremendous amount of paper that required duplicated file storage and maintenance. By replacing the paper personnel file with an electronic file, the County has created a comprehensive electronic personnel data storehouse that spans the lifecycle of the employee.

Using integrated content and records management systems, Buncombe County shares information between departments, has reduced storage requirements, and has drastically improved efficiencies. There is no more duplication, the amount of work has been greatly reduced, and as a rule, documents do not get lost and departments have the right information. The County is also provided with a complete audit trail of all transactions and personnel who have access to official records.

Buncombe County's success proves that the path to agentic intelligence begins with a records-first mindset. By eliminating duplicate files and establishing a complete, auditable data lineage, they have done more than just save office space—they have built the enterprise memory required for agents to act with confidence.

Content Systems: The Memory Layer of the Agentic Genome

To support a dynamic, living data ecosystem, enterprises must view their information architecture differently. Enterprise content systems include:

- Knowledge bases (policies, guides, standard operating procedures)
- Document repositories (PDFs, word processing docs, spreadsheets, reports)
- Data lakes and warehouses (raw or aggregated data, logs, usage data)
- Operational databases (CRM, ERP, billing systems, order systems, user databases)
- Supply chain or business transactional data (real-time and data repositories)
- APIs and microservices (internal/external) that expose data or business logic
- Legacy systems (mainframes, older databases, and file shares)
- Metadata, audit logs, and identity and access management systems

Together, these content systems and the data they manage provide the “memory” of enterprise information—the input, context, historical state, and infrastructure information—that agents need to perform complex workflows.

As discussed in the first book in this series, every organization generates a vast amount of information for a variety of purposes, including documenting business operations, communicating with internal and external stakeholders, preserving knowledge, meeting regulatory obligations, and ultimately, delivering value to customers and partners.

What started as structured management of content, including transactions, invoices, and documents, has steadily expanded into a sophisticated mesh of communication, collaboration, and digital interactions—from emails, shared documents, and messaging platforms to workflows and case files.

Today, this information can be understood as three primary classes of enterprise data, each shaping AI in different ways: 1) human-generated content, 2) machine-generated content, and 3) transactional or business network data. Each category necessitates its own structure and governance requirements. Together, they form the foundation of agentic intelligence inside the modern enterprise.

Human-Generated Content: The Intent Layer

This class of content encompasses documents, messages, scanned records, multimedia, and every other form of human communication. Inherently unstructured, this content spans the breadth of the enterprise. It often carries personal, contextual, and confidential elements that require classification, metadata enrichment, access controls, and lifecycle governance.

Human-generated content also includes the organization's rules, policies, precedents, procedures, and professional judgment. Allowing AI systems to leverage this material requires de-identification, accurate metadata, and strong governance. This is where intent, business context, and organizational expertise reside.

When curated correctly, unstructured content can serve as the foundation for RAG pipelines, prompt libraries, generative reasoning, conversational agents, and domain-specific copilots, enabling AI agents that move beyond task automation to truly understand the business they operate in.

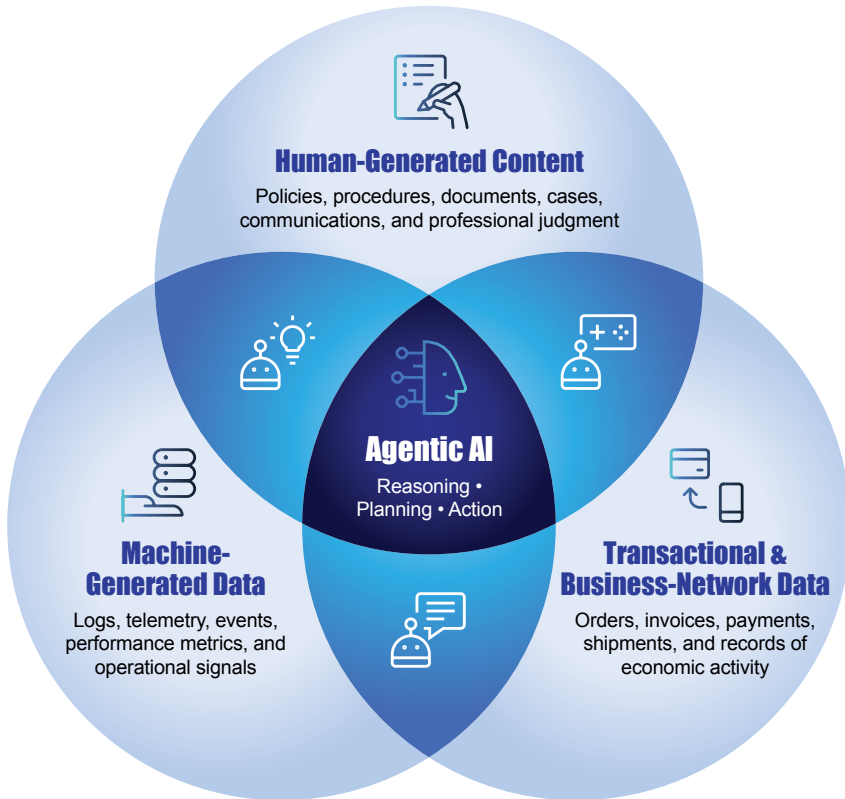
Machine-Generated Data: The Operational Nervous System

Machine-generated data—including logs, telemetry, performance metrics, events, and monitoring streams—comes from the systems that run the organization. High-volume, high-velocity, and highly structured, this data tells the enterprise story in real time. It represents the practice of observability and operational intelligence, where every event, anomaly, or deviation leaves a measurable trace.

Machine data provides the signals that allow AI to act with operational awareness: detecting patterns, diagnosing failures, predicting outages, optimizing workloads, or initiating corrective workflows before users are impacted. Its challenge lies in managing scale, controlling retention costs, and transforming raw signals into a business context.

When combined with human-generated content and policy frameworks, machine-generated data enables agentic systems to act autonomously while remaining traceable, governed, and aligned with enterprise expectations.

Three Primary Classes of Enterprise Data Shaping Agentic AI



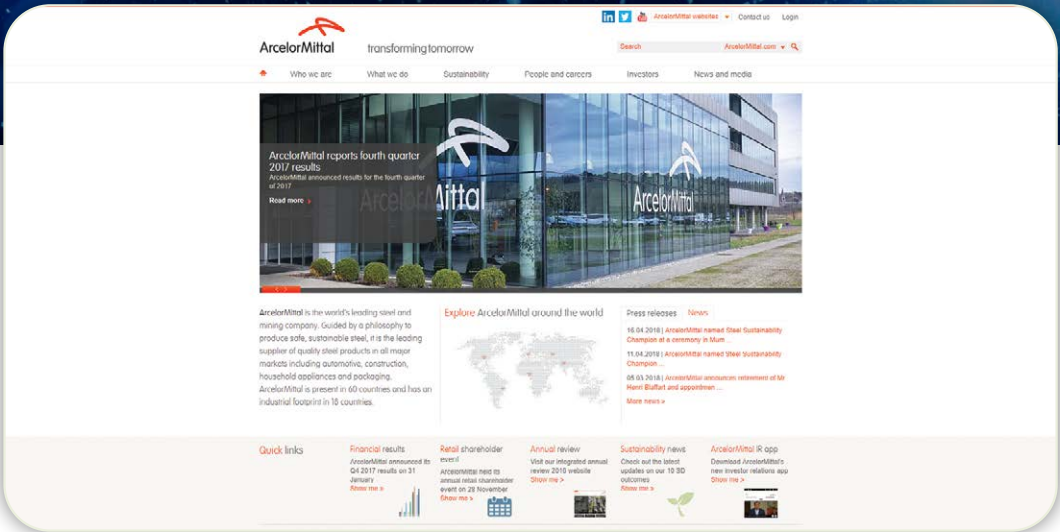
Each Data Class Shapes How Agents Reason, Observe, and Act

Transactional and Business Network Data: The Enterprise Source of Truth

Transactional data, including purchase orders, invoices, shipping updates, account records, and other structured business messages, represents the legal and operational truth of the enterprise. These records define obligations, financial positions, compliance requirements, and the flow of economic value across business networks.

Because transactional systems follow strict schemas and carry high semantic precision, they provide a foundation of trust for reasoning. For agentic AI, this is the anchor to reality. It enables agents to reconcile financial positions, forecast cash flow, detect anomalies, process claims, validate compliance, and monitor supply chain performance with confidence.

While human-generated content provides intent, it is the transactional layer—the single source of truth—that anchors an agent's reasoning in reality. ArcelorMittal demonstrates this by transforming a fragmented patchwork of 100,000 monthly messages into a governed, high-visibility business network that serves as the essential bedrock for global supply chain orchestration.



ArcelorMittal

ArcelorMittal is the world's leading steel and mining company. Guided by a philosophy to produce safe, sustainable steel, it is the leading supplier of quality steel products in all major markets including automotive, construction, household appliances, and packaging. ArcelorMittal is present in 60 countries, has an industrial footprint in 18 countries, and ranks as a truly global steelmaker.

The company wanted to accelerate the process of onboarding new partners to its B2B e-commerce network while enhancing their capability to deliver new solutions. In addition to being time consuming, running B2B applications in-house was proving to be costly. As such, the company also sought to reduce the support and maintenance costs of older mainframe-based B2B applications that had been developed in-house. Outsourcing a patchwork of fragmented B2B trading networks via B2B Managed Services emerged as an effective and powerful way to address all of these challenges.

Entrusting B2B Managed Services with the management of more than 100,000 monthly Electronic Data Interchange (EDI) messages between global trading partners has eliminated the need for in-house B2B technology expertise. It has resulted in smoother and more reliable core business processes, increased global visibility across the business network, and improved monitoring capabilities. Based on the initial success ArcelorMittal has experienced, the company plans to extend the reach of B2B Managed Services to include additional trading partners.

A global footprint is very important to us as we need to be able to connect to a customer anywhere in the world. B2B Managed Services increases global visibility across our base of B2B transactions and delivers considerably improved monitoring capabilities.

ArcelorMittal have created the transactional memory necessary for a global enterprise to act as a single, cohesive unit. This foundation does more than just reduce IT maintenance; it provides the semantic precision required for future agents to reconcile global shipping, manage partner obligations, and optimize the flow of sustainable steel at machine speed.

Content Systems in Concert

When transactional accuracy is combined with the context of human-generated knowledge and the signals of machine-generated telemetry, enterprises gain a complete, dynamic picture of their operations: what is happening, why it is happening, and what the system should do next.

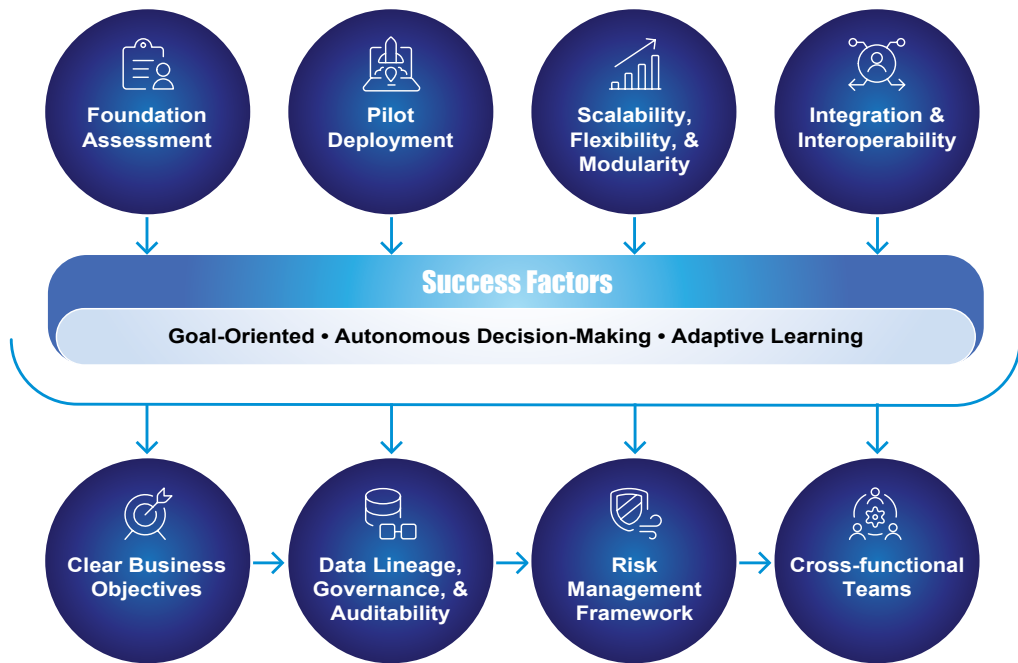
Deep understanding of this content is vital to drive agentic behavior, allowing agents to reason within real-world business contexts. Further, agents often need cross-system visibility (for example, reading CRM records, checking the billing system, and consulting compliance policies) to perform complex tasks.

Crucially, these systems also serve as gatekeepers, defining exactly what agents can—and cannot—access. Access rights, permission boundaries, and audit logs are all key to establishing enterprise-grade trust, compliance, and security.

These content systems are the backbone of an agentic AI architecture.

Architectural Requirements for an Agentic Content Foundation

Given the critical role of content systems in enabling agentic AI, specific architectural imperatives emerge. Without them, agents will lack the capability to work across domains and dynamically adapt to change.



Strategic Agentic AI Implementation Framework

If content systems are the memory layer of the agentic enterprise, then their architecture determines whether they can function at scale. Three pillars are essential:

1. Integration and Interoperability

Content systems must be easy to integrate with, providing APIs, standard protocols, and well-defined metadata. Agents require consistent access to the full spectrum of enterprise information. Consequently, indexing, search, and retrieval mechanisms must equally support queries and data that are structured, semi-structured, and unstructured.

Metadata and indexing layers (e.g., full-text, vector, and semantic search) should be implemented to enable agents to locate relevant content efficiently.

2. Data Lineage, Governance, and Auditability

Every piece of content must carry provenance: origin, ownership, version history, permissions, and policy constraints. In agentic systems, auditability must extend to actions, not just reads. This is the basis of defensibility in regulated and sovereign contexts.

Governance policies must be strictly enforced, including who can read/write, consent for sensitive data use, retention policies, and compliance with internal/external regulations. Audit trails and logging are paramount, especially when agents execute tasks (not just read), to ensure traceability and accountability for decisions, updates, and actions.

1. Scalability, Flexibility, and Modularity

As the enterprise evolves with new business units, data sources, and content types, the content architecture must support incremental growth without creating technical debt.

Adopting a modular design ensures each content system (or domain) can evolve independently while remaining accessible via standard protocols/interfaces. Furthermore, architectures must feature flexible schema and metadata handling to support new data types (e.g., multimedia, sensor data, streaming logs) as business needs evolve.

These requirements ensure that agents have a robust, scalable, maintainable, and secure foundation—one that can keep up with business change while preserving trust and control. Without these properties, agentic systems remain confined to pilots. With them, they become platforms.

Designing Content Systems for Agentic AI: Principles and Best Practices

In an agentic enterprise, the cloud is not simply a hosting environment; it is the enabling fabric for intelligent execution at scale. Agentic AI depends on scalable processing power (elastic compute), real-time data access, secure integration across systems, and continuous orchestration of workflows that span organizational and geographic boundaries.

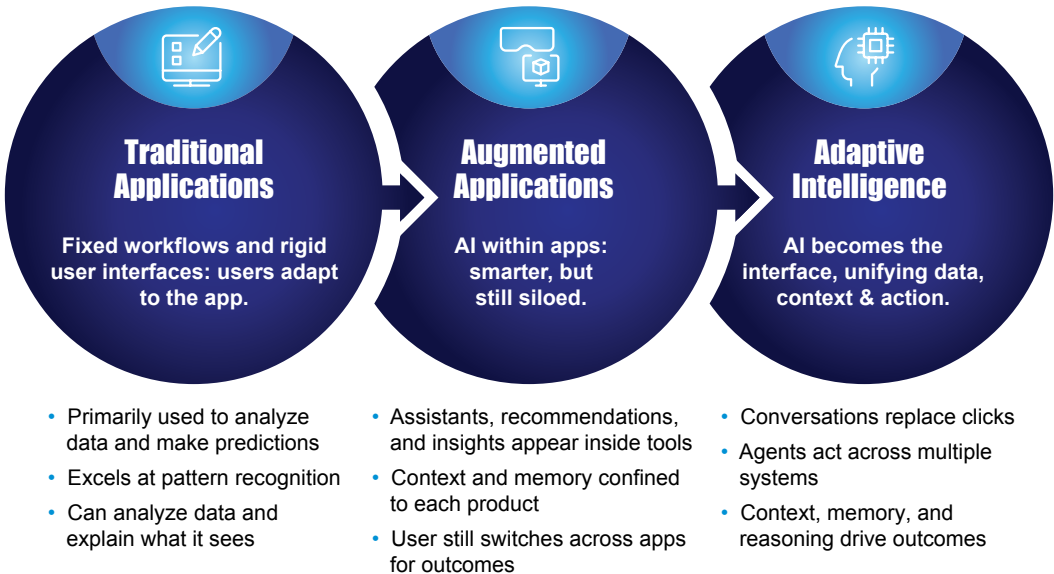
Cloud platforms provide the scalability, resilience, and connectivity required for agents to operate continuously, coordinate across domains, and adapt as demand fluctuates. Crucially, they support governance-by-design—identity, access control, auditability, and policy enforcement embedded into the platform itself—allowing agentic systems to scale without sacrificing trust, compliance, or control.

However, designing for an agentic enterprise requires more than just cloud infrastructure; it requires an architectural approach that ensures accuracy, trust, interoperability, and control.

Single Sources of Truth

The first core principle is the source-of-truth model. Organizations must avoid the proliferation of redundant or inconsistent data, which can introduce misleading AI outcomes that undermine trust. Instead, systems of record, such as CRM platforms, ERP systems, and compliance databases, should serve as the sources from which agents read. By anchoring agentic operations to verified, governed data stores, enterprises can reduce errors and prevent automated actions from amplifying errors across downstream systems.

Evolution: Static applications to adaptive intelligence



The Agentic Enterprise UI Will Be Powered by Context and Intelligent Actions

Context By Design

A second essential requirement is the deliberate enrichment of content with metadata and contextual layering. Data must carry more than its raw value; it must include the information that gives it meaning. Each object, whether a document, record, or event, should carry attributes such as timestamps, ownership details, version history, quality indicators, access permissions, and semantic context tags.

These metadata layers allow both human users and AI agents to assess whether the content is valid, current, trustworthy, and appropriate for a given task. In an environment where autonomous systems act on information, context is as essential as the content itself. Context determines whether an action is appropriate, permitted, or defensible.

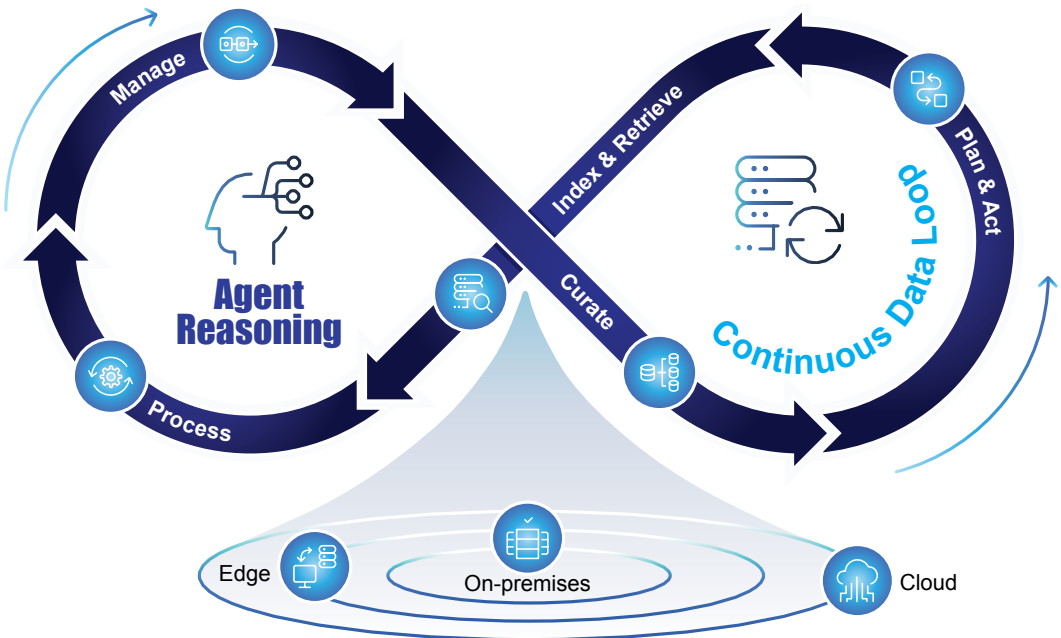
API-First, Tool-Ready Platforms

Equally important is adopting an API-first, “tool-ready” design philosophy, where content systems are exposed as reliable services rather than static silos. In other words, APIs and software interfaces should be designed specifically so that AI agents can “read” the documentation, understand what the tool does, and execute it autonomously. Systems need to provide well-documented APIs, searchable indexes, and event-driven services that allow agents to retrieve content, request state changes, initiate workflows, and subscribe to updates. This shift eliminates legacy workarounds and replaces them with structured, governed interfaces that support automation at scale. By treating content as a service—and software features as “tools” that agents can wield—enterprises create a flexible platform that agents can interact with safely and predictably.

Governance

Governance binds these principles together, integrating security, access control, compliance enforcement, and auditability. Because Agentic AI may read, write, and execute complex tasks, governance cannot be an afterthought. Least-privilege access (where users only have access to what is absolutely necessary for them to perform their roles), robust identity management, permission boundaries, continuous logging, and escalation mechanisms ensure that agent behavior remains observable, reversible, and aligned with enterprise policy. This discipline preserves trust, especially as autonomous workflows become more frequent and complex.

When enterprises adopt these principles holistically, they create a content foundation that enables agents to operate confidently and responsibly. The result is an architecture that preserves control and compliance while unlocking the adaptability and efficiency required for agentic AI.




Agentic AI Continuous Data Loop

When the three pillars of an agentic foundation—interoperability, lineage, and scalability—are applied to the enterprise’s most sensitive data, the architecture shifts from a passive storage layer into an active defense mechanism. The following case study of a global high-tech manufacturer illustrates how this operational nervous system can transform raw security telemetry into autonomous, goal-oriented threat hunting.



A High-Tech Manufacturer



A global high-tech manufacturer faced a growing and sophisticated threat landscape: with tens of thousands of employees and valuable intellectual property at stake, traditional security monitoring tools were struggling to keep up with targeted external attacks and insider threat risks. Rather than relying solely on reactive alerts and human-centric investigation workflows, the organization could adopt a more autonomous, agentic approach to cybersecurity analytics.

In an agentic deployment, unsupervised learning models serve as security reasoning agents that continuously ingest and interpret a wide variety of operational signals—from data loss prevention (DLP) endpoint telemetry and Active Directory events to email patterns and source code repository logs. These analytic agents detect subtle behavioral shifts and contextually relevant risk indicators that would otherwise remain hidden. For instance, an event trigger such as unusual access to source code automatically elevates a case into an adaptive threat-hunting workflow, providing the context by design required to distinguish a legitimate developer's work from a potential data exfiltration attempt.

An agentic architecture transforms security operations from rules-based alerts to continuous, goal-oriented insight. Instead of waiting for predefined thresholds to be crossed, analytic agents proactively reason about the context of activity patterns, refine their models over time, and drive dynamic prioritization for human defenders. The result is a more resilient insider threat program that preempts attacks, reduces analyst churn, and aligns detection with evolving enterprise risk—no matter how fast the threat environment changes. In an era of machine-speed attacks, this gives the agentic enterprise a self-shielding architecture that is as adaptive as the threats it faces.

Implications: What This Means for Enterprise Leaders and Architects

The shift from static AI inputs to living content ecosystems and agentic AI reframes enterprise strategy. It demands a fundamental rethink of the organization's digital foundation.

Leaders must ask: *Is our information architecture agent-ready? Can it support real-time reasoning, cross-domain orchestration, and governed action?*

Leaders and architects must begin by reassessing the current state of their information landscape. Many organizations still rely on siloed data stores, file shares, and legacy systems that lack APIs, versioning, metadata enrichment, or modern access controls.

The critical first step is to identify gaps. In many cases, this evaluation will reveal that current systems cannot support agents that depend on real-time data, contextual awareness, and governed access pathways. These gaps must be addressed before agentic AI can be deployed effectively.

To meet these new requirements, enterprises must invest in "agent-ready" content foundations—systems that are modular, API-first, metadata-rich, and designed for continuous evolution. This means modernizing legacy databases, integrating systems through standards-based interfaces, and ensuring that content is not merely stored but is discoverable, contextualized, and actionable. Such systems will form the operational foundation for agents, enabling them to read, reason, update, and collaborate across domains without compromising the integrity or governance of enterprise information.

Secondly, enterprises should embrace hybrid AI architectures that balance retrieval-based intelligence with agentic capabilities. Retrieval-Augmented Generation (RAG) remains invaluable for grounding language models in enterprise knowledge, enabling accurate and contextually relevant responses.

However, when the business requires orchestration, decision-making, workflow execution, or multi-step reasoning, agentic architectures are essential.

Rather than replacing one another, RAG and agentic AI should be viewed as complementary components within a unified ecosystem. The organizations that succeed in this transition will be those that architect their content systems to support both, delivering intelligence that is not only informed but actionable.

At the same time, leaders must broaden their understanding of data governance and compliance. Protecting data itself is no longer sufficient; organizations must now safeguard agent behaviors. Actions taken autonomously—whether updating a record, triggering a workflow, or making a recommendation—must be auditable, traceable, and aligned with internal policies and regulatory frameworks. Building governance into the architecture from the start ensures accountability and reduces operational risk.

Executive Implications

CAIO

Agentic AI requires governance by design. The CAIO establishes the policies, access controls, and audit frameworks that ensure enterprise information can be safely used by agents while maintaining sovereignty, accountability, and compliance.

CIO

Agentic AI will not scale on models alone; it requires an agent-ready content architecture—API-first, interoperable, and modular—so agents can read, write, and act safely across enterprise systems.

CFO

Information architecture becomes a productivity and risk lever: investing in single sources of truth, metadata, and auditability prevents costly automation errors and turns AI spend into scalable operating capability.

CHRO

As policy and people data become operational inputs for agents, HR must strengthen access controls, lifecycle governance, and human-in-command practices to protect sensitive employee information while enabling reliable self-service.

CDO

Data governance must expand from protecting information to governing behavior, ensuring lineage, context, permissions, and audit trails are embedded so autonomous actions remain defensible under scrutiny.

COO

Operational speed and resilience depend on content systems working in concert—combining human knowledge, machine telemetry, and transactional truth—so orchestrated workflows can execute in real time without sacrificing control.

In the next chapter, we'll examine how sovereign cloud intersects with agentic AI and why agentic AI demands governance, human-in-command design, and defensibility.



Chapter Three

Sovereignty, Accountability, and Control

Chapter 3 is about the conditions that make agentic AI safe, scalable, and defensible in the enterprise. As AI systems move from generating content to executing work, questions of sovereignty, accountability, and control move from architectural considerations to executive responsibilities. This chapter examines why agentic AI must be grounded in sovereign deployment models, governed through human-in-command oversight, and designed with auditability and policy enforcement built into execution—not bolted on after the fact.

Artificial intelligence has evolved from being just an analytical tool for search and summarization to becoming an integrated part of the enterprise workforce. As the capability of AI increases, so too do the potential consequences—both positive and negative. Agentic AI demands a deliberate architecture of control and, in some cases, sovereign control of the technical, operational, and legal aspects of AI. The most successful organizations of the next decade will not be those who build the largest models, but those who govern and deploy AI that people can trust.

Why Agentic AI Demands Governance

Traditional AI systems primarily generate insights. Agentic AI systems generate actions.

When AI agents can update records, trigger payments, approve permits, optimize logistics, or triage cases, the risk profile shifts from a simple informational error to a more serious operational consequence. The governance challenge, therefore, expands beyond model accuracy to include:

- Decision authority
- Execution control
- Auditability
- Compliance
- Legal and Regulatory Defensibility

As described in Chapter 1, data is the foundational asset that fuels these systems. It drives innovation, productivity, and competitive advantage. But as sovereign AI frameworks emphasize, data is also a matter of national and economic security. AI's powers of inference mean that even small pieces of data can produce sensitive insights when aggregated. In the agentic era, every piece of data has the potential to become consequential. That means strong governance is crucial.

Governance for agentic AI should extend across three domains:

1. **Data Governance** – ensuring data quality, classification, access control, and retention.
2. **Model Governance** – validating models and ensuring explainability and fairness.
3. **Action Governance** – controlling how and when agents act, including escalation pathways and human oversight.

The third domain, action governance, is what distinguishes agentic AI from earlier implementations of enterprise AI.

Human-in-Command Design

As AI systems become more autonomous, human responsibility must become more prevalent.

Human-in-command design means that humans retain ultimate authority over the boundaries of agent action, escalation thresholds, ethical interpretation, and strategic decision making.

This does not mean that humans manually approve every action—as we pointed out in Chapter 1, this is where the human-in-the-loop can create bottlenecks. Rather, organizations must deliberately define autonomy tiers for agentic AI. For example:

- **Tier 1: Advisory agents** (recommendation only, acting as an assistant to the human in the role)
- **Tier 2: Assisted execution** (agent acts with audit logging and reversible workflows, with human oversight)
- **Tier 3: Basic autonomy** (agent acts within defined thresholds and for controlled workflows, with human auditability)
- **Tier 4: Full autonomy**

Human-in-command design also requires increased mechanisms for transparency, including the full logging of agent reasoning paths, audit trails linking data sources to decisions, explainability frameworks, and override and rollback capabilities.

In this case, humans define risk trade-offs. They resolve ambiguous ethical dilemmas and they interpret context beyond structured inputs.

Agentic AI must be treated not as a replacement for human domain expertise, but as a force multiplier, capable of driving enhanced productivity and outcomes while preserving judgment and accountability.

To illustrate, let's take a look at how the enterprise operations role of a Site Reliability Engineer (SRE) is impacted through the introduction of agentic AI capabilities, to help ensure the performance of critical IT systems and applications across an organization.

Agentic AI and the Site Reliability Engineer (SRE)



SRE Agentic AI “Assistant”

Scope of Work:

- Infrastructure health (Kubernetes, cloud platforms, middleware)
- Reliability metrics (SLOs, error budgets, availability)
- Platform troubleshooting and diagnostics
- Container orchestration and deployment health

Key Activities:

- Infrastructure triage and blast radius assessment
- Troubleshooting guidance and diagnostic recommendations
- Runbook retrieval for known patterns
- Root cause analysis support
- Health assessments on demand
- Post-incident timeline reconstruction



SRE Human Role

Scope of Work:

- Executes remediation commands
- Restarts pods, nodes, or services
- Modifies infrastructure configuration
- Makes scaling decisions
- Defines architecture

In this example, the SRE Agentic AI Assistant focuses on sensing, analysis, and decision support across the operational environment. It continuously monitors infrastructure health and reliability signals, correlates telemetry across platforms, retrieves relevant runbooks, and provides diagnostic guidance and post-incident analysis. In other words, the agent specializes in perception, reasoning, and situational awareness at machine speed, helping surface the right information and recommended actions at the right time.

The human SRE, by contrast, retains authority over execution and architectural judgment. Humans carry out remediation actions, make scaling and configuration decisions, and shape the underlying system design. This division preserves accountability and risk management: agents accelerate insight and response, while humans remain responsible for changes that affect system behavior, resilience, and long-term architectural integrity.

Defensibility of Agentic AI

Defensibility is the ability to explain, justify, and reconstruct decisions. This is important in the context of agentic AI, where humans are not always making decisions. Further, in the public sector and in regulated sectors such as banking, healthcare, and telecommunications, actions must withstand scrutiny from auditors, regulators, courts, and the public. This requires enhanced defensibility measures, including but not limited to:

- **Immutable logs** – digital records of system activities, events, or data that cannot be altered or deleted
- **Data lineage tracking** – automated mapping of the entire lifecycle of data
- **Version-controlled model artifacts** – saved outputs of the machine learning process
- **Clear policy mapping** – visualizing, defining, and organizing data to highlight where policy intervention is required
- **Jurisdiction-aware data handling** – identifying, tracking, and managing data based on the specific legal, regulatory, and geographical boundaries governing it

Trusted data and responsible AI are inseparable. If the data is poorly governed, poorly contextualized, and not sovereign, the AI built upon it cannot be defensible. Organizations must ask the following basic questions to ensure that their defensibility strategy is comprehensive:

Who owns the data?

Where does it reside?

Who has access?

Under which laws is it governed?

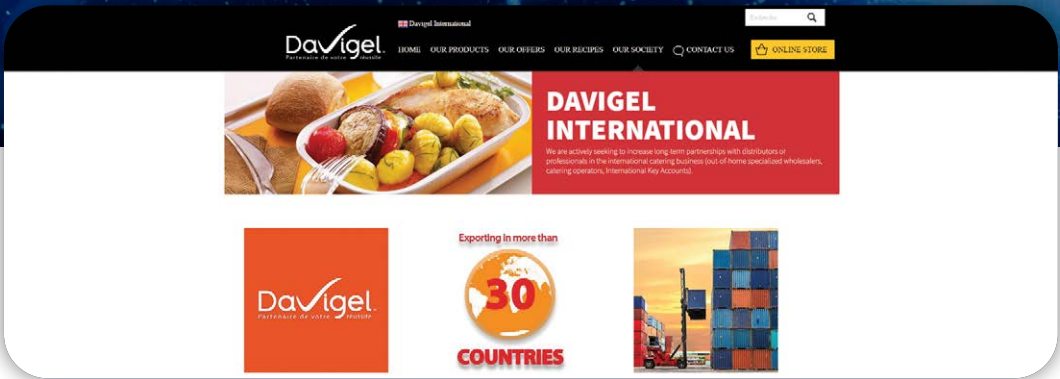
How long is it retained?

Agentic AI multiplies the importance of these questions. Decisions made by agents become institutional decisions. Institutional decisions carry legal consequences.

Defensibility is therefore not a compliance afterthought. It is a core design requirement.

The transition to agentic commerce requires governance that can navigate the shifting sands of global regulation. In the following case study, Davigel moves from manual, high-risk billing to a sovereign e-invoicing engine that automatically enforces the distinct legal and tax mandates of over 45 jurisdictions.

Davigel



Davigel International

Davigel SAS, owned by SYSCO, is a leading producer, importer, and distributor of branded frozen and chilled food products across the commercial, social, and airline catering sectors. The company provides an innovative range of products and solutions for more than 65,000 customers around the world.

Shortly after being acquired, Davigel partnered with a Managed Services provider to migrate its trading partner community of customers and logistics providers to its own B2B trading platform. Building on the success of the B2B migration, Davigel expanded its Managed Services relationship to include an integrated e-billing solution. A key factor in electronic billing is ensuring that each invoice is compliant with local regulations. Every country in the world has legislation to regulate electronic invoicing, and while regulations are often similar in purpose, the specific requirements can vary significantly by country. For example, data archiving requirements can vary from 6 to 11 years depending on which countries are involved. In addition, having to manage the country-specific value-added tax (VAT), security, and archiving requirements for every customer significantly complicated the billing process.

To address this challenge, Davigel deployed a secure, cloud-based e-billing solution with the capability to automate invoicing and compliance in more than 45 countries. Compatibility with the company's recently migrated B2B trading platform enabled Davigel to establish an e-invoicing program that could support data validation, data signatures, archiving, and delivery—all while providing employees and auditors with access to the archive via a simple, intuitive web portal. Thanks to the e-invoicing program, Davigel can now confidently issue over 120,000 invoices to more than 65,000 global customers securely and in accordance with regional regulations.

Sovereignty is not just about where data sits, it's also about how laws are encoded into action. Davigel's model proves that when accountability is built into the execution layer, even the most complex global workflows can move at the speed of the cloud without sacrificing the human-in-command oversight necessary to satisfy auditors and regulators alike.

Sovereign Cloud and Agentic AI

Sovereignty becomes central when AI systems operate at scale. Digital sovereignty refers to a nation or organization's ability to maintain control over its digital assets, systems, and operations, independent of external jurisdictional influence and in compliance with domestic law.

Agentic AI intensifies sovereignty concerns because:

- Agents may act across borders
- Data may move dynamically between systems
- Control planes may reside in foreign jurisdictions
- Model training and inference may involve sensitive datasets

Four elements of sovereignty must be considered in the context of agentic AI:

1. Data Sovereignty

Sensitive data must be stored, processed, and managed within defined jurisdictions. Agents operating on sovereign data cannot expose it to foreign legal claims or external cloud providers without appropriate safeguards.

2. Operational Sovereignty

Employees managing AI systems must operate within trusted jurisdictions. Clearance levels, identity verification, and operational boundaries matter especially in the public sector and regulated industries.

3. Technological Sovereignty

Control over infrastructure, encryption keys, hardware, software, and control planes must remain under sovereign authority. Without technological sovereignty, governance policies can be bypassed at the infrastructure layer.

4. Legal Sovereignty

Cloud providers and AI vendors must be governed under appropriate legal frameworks to prevent legal exposure.

While the theoretical framework of sovereignty defines the boundaries of control, the operational reality for a global engineering and infrastructure enterprise requires an architecture capable of enforcing those principles across widely distributed environments. Organizations operating across dozens of countries and remote industrial sites must harmonize the four pillars of sovereignty—data, operational, technological, and legal—while ensuring that critical information remains secure, accessible, and governed wherever work occurs. This is demonstrated in the following case study.



A Global Engineering and Infrastructure Firm



A global engineering and construction company operating major energy and infrastructure projects across offshore and onshore environments employs tens of thousands of workers and manages operations in more than fifty countries. Many of its projects take place in remote industrial locations—including offshore platforms and field construction sites—where connectivity can be intermittent. Engineering drawings, project documentation, safety records, and contractual agreements form the operational backbone of these projects, making secure and compliant information management essential to the lifecycle of every asset.

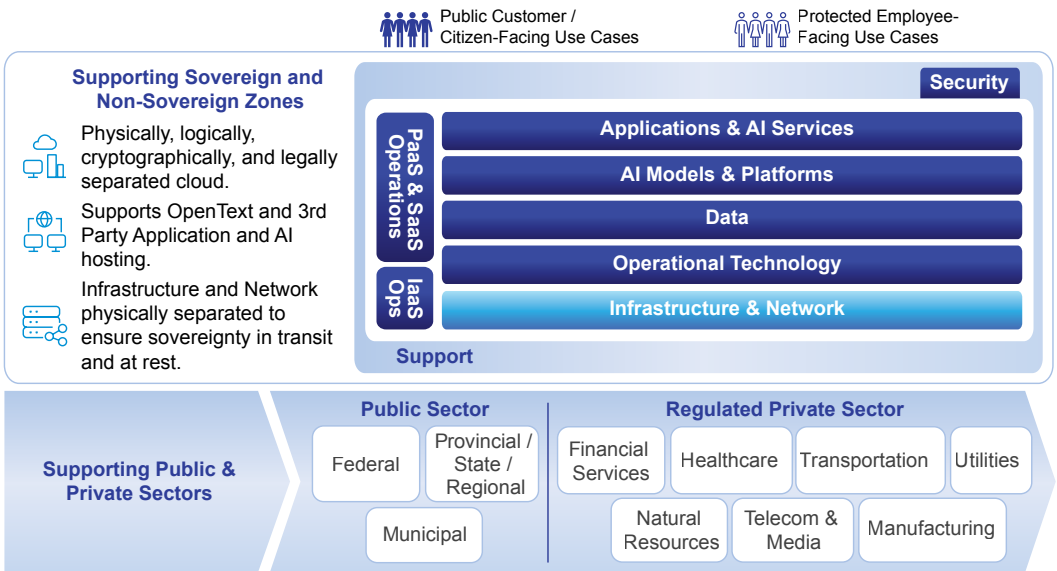
As part of a broader move toward cloud-based platforms and as-a-service infrastructure, the organization migrated its enterprise content management environment from on-premises systems to a cloud architecture. The transition required careful validation of workflows and the migration of more than 60 terabytes of historical engineering and project data. Once completed, the platform delivered measurable operational improvements: end-user performance increased by more than 50%, technical support requests fell by roughly 90%, and operational costs declined by approximately 30%.

Beyond efficiency gains, the cloud-based content platform created the “sovereign memory” required for intelligent automation and AI-driven information workflows. With governed enterprise content accessible through secure APIs, intelligent systems can begin to analyze engineering documents, surface operational insights, and support automated workflows across distributed project teams. In this environment, AI is not simply an analytics tool—it becomes part of a governed information architecture that allows global engineering operations to scale securely while maintaining sovereignty, compliance, and operational control.

The Hybrid Sovereign Model for Agentic AI

To compete in the AI era, enterprises and nations must leverage the scale and innovation of global hyperscalers. Yet full reliance on external providers introduces risk. The solution is a **balanced hybrid model**.

In the first book in this series, *Enterprise Artificial Intelligence: Building Trusted AI in the Sovereign Cloud*, we introduced a hybrid architecture for sovereign data and AI.



High Level Architecture of Sovereign Data and AI

This architecture establishes:

- A Private Sovereign Zone for sensitive datasets and high-trust workloads.
- A Public or Hybrid Zone for non-sensitive, citizen-facing, or scale-intensive workloads.

Within this model:

- Private AI Agents operate inside the sovereign stack, analyzing and acting on protected data without risk of leakage.
- Public AI Agents operate in scalable cloud environments, delivering digital services and interacting with less sensitive datasets.

This dual-zone model enables both control and competitiveness.

The infrastructure layer employs zero-trust protocols, air-gapped configurations where necessary, and security-cleared operational oversight. The operational technology layer standardizes governance across multi-cloud environments through unified control planes and orchestration tools. The data layer integrates explicit hierarchies (schemas, taxonomies, version control) with implicit semantic structures (metadata, ontologies, usage clustering), enabling AI systems to reason across complex information domains while respecting policy boundaries.

This architecture provides the secure foundation required for agentic deployment.

Sovereignty as Strategic Advantage

In a world of geopolitical challenges, increasingly complex regulatory frameworks, and fear of AI influence, sovereignty is not a defensive position; it is a competitive strategy.

Organizations that build jurisdiction-aware, resilient, and policy-aligned AI architectures gain:

- Operational continuity under disruption
- Legal certainty
- Public trust
- Strategic flexibility

Data is not a commodity—it is a competitive differentiator for public and private-sector enterprises. It is something to be protected and deployed responsibly. The enterprises and countries that will lead in the cognitive computing era will not simply build powerful agents. They will build *governed agents* that are embedded in architectures that align innovation with accountability.

When innovation and governance are aligned and when privacy, accountability, sovereignty, and ethics are embedded by design, agentic AI becomes not just intelligent, but trustworthy.

Executive Implications

CAIO

Agentic systems require governance as architecture. The CAIO defines autonomy boundaries, risk classification, and cross-functional oversight so human-in-command design is enforceable at enterprise scale.

CIO

As AI moves from insight to action, enterprise architecture must embed sovereignty, auditability, and control planes by design so agentic execution can scale without introducing systemic risk.

CFO

Autonomous execution changes the enterprise risk profile; defensible ROI now depends on governance, traceability, and jurisdiction-aware deployment that prevent costly compliance failures and operational exposure.

CHRO

Human-in-command is a workforce design choice: roles, escalation paths, and accountability must be redefined so people supervise digital labor rather than perform every operational step themselves.

CDO

Data governance must expand into action governance, ensuring lineage, access, explainability, and policy mapping so agent-driven decisions remain transparent and legally defensible.

COO

Operational resilience requires tiered autonomy and orchestration that allow agents to accelerate routine execution while preserving human authority over high-impact decisions and systemic change.

In the next chapter, we explore how to redesign operating models for human–AI collaboration, including the orchestration of complex, end-to-end workflows with multiple agents across the enterprise.

Chapter Four

From Org Charts to Orchestration

In this chapter, we'll discuss the shift from AI as a personal productivity layer to AI as a coordinated operating capability. We explore why accelerating individual tasks does not transform enterprise outcomes, and how nested orchestration enables intelligence to participate directly in workflows with accountability and control. The result is a new cognitive operating model that replaces functional silos with orchestrated flows of information, decisions, and action.

Redesigning the Operating Model for the Cognitive Enterprise

In the previous chapters, we explored how agentic AI transforms enterprise systems from sources of insight into engines of execution. As intelligence begins to participate directly in workflows, the question shifts from technological capability to organizational design. The challenge is no longer simply deploying AI tools; it is redesigning how work flows through the enterprise.

This chapter examines that transition.

Early enterprise deployments of generative AI focused on personal productivity. Virtual assistants summarized emails, drafted proposals, generated code, and accelerated documentation. AI appeared as a digital assistant layered onto the daily work of knowledge professionals. The logic was simple. If every employee could think and write faster, the enterprise would move faster.

And for a moment, it appeared to work.

Tasks were completed more quickly. Drafts improved. Friction at the level of the individual decreased. Teams felt empowered. Yet across industries, a more sobering reality began to surface. Despite broad experimentation and significant investment, few organizations experienced the structural performance gains that boards expected. Costs rose. Architectural complexity increased. Governance risks multiplied. The anticipated step change in enterprise performance remained elusive.

The issue was not model intelligence; it was operating model design.

Enterprise value is rarely created by isolated individuals working faster. It is created through the coordinated flow of information, decisions, and actions across interconnected processes. A sales team may draft proposals more quickly, but if credit approvals stall in finance, revenue still slows.

Speed at the edge cannot compensate for friction in the core.

When AI is treated primarily as a productivity layer, it accelerates activity without transforming outcomes. The enterprise becomes busier, not necessarily better. For the C-Suite, this moment represents an architectural inflection point. The question is no longer: *where can AI assist employees?* The deeper question is how to architect the enterprise so that intelligence becomes embedded directly into execution.



From Org Charts to Orchestration with Agentic AI

Previous transformations digitized workflows. The agentic transformation digitizes decision-making itself. Agentic systems do not simply generate outputs. Properly architected, they perceive context, maintain state, plan actions, make bounded decisions, execute work, and manage exceptions within defined governance constraints.

Once intelligence begins to participate directly in workflows, the operating model must evolve to accommodate it.

The Case for Structured Intelligence

Early AI deployments often assume that intelligence can be centralized into a single, powerful system. The vision of an enterprise-wide super system managing everything from procurement to compliance is attractive. It promises simplification and unified insight.

In practice, it does not scale.

Enterprises are complex because their environments are complex. Regulatory requirements vary by geography. Data access is bounded by policy and law. Functional domains require specialized reasoning. Risk tolerances differ across processes. Decision rights are distributed for good reason.

A monolithic AI system attempting to manage all enterprise activity becomes brittle. It struggles with explainability, introduces security exposure, and becomes difficult to audit. Most importantly, it erodes trust.

Human organizations have evolved their hierarchy as a mechanism to manage complexity. Hierarchy distributes responsibility, limits span of control, and embeds governance into structure. Agentic AI requires the same discipline.

What emerges isn't centralized intelligence, but structured intelligence organized in nested layers of orchestration.

At the lowest level, specialized agents perform narrow tasks with precision. Above them, operational orchestrators coordinate workflows. Higher still, domain orchestrators enforce policy and regulatory constraints. At the top, enterprise orchestration aligns activity across value streams and strategic objectives.

Each layer reduces complexity for the layer above it while enforcing governance for the layer below it.

This layered structure allows intelligence to scale without sacrificing control. Instead of a single system attempting to manage every decision, the enterprise distributes agency across orchestrated tiers that mirror how organizations already manage complexity, risk, and accountability.

This model—nested orchestration—forms the cognitive architecture of the agentic enterprise.

The Modern Org Chart – Encompassing Your Human and Digital Workforce

At the operational level, organizational structures must reflect the hybrid architectural reality of modern AI. This means defining roles that manage both data infrastructure (including quality, lineage, and access), and action governance (including decision frameworks, audit trails, and human oversight gates). Architects, data engineers, and business analysts should work in integrated squads rather than traditional, siloed departments. Cross-functional teams allow technical, operational, and compliance considerations to be addressed simultaneously.

Equally important is workforce enablement. Employees must understand how to collaborate with agents, interpret their outputs, and intervene when necessary. Training programs, clear escalation pathways, and defined accountability foster confidence rather than resistance.

Even as autonomous agents assume greater operational responsibility, humans remain central to enterprise effectiveness. Humans define the boundaries of agent autonomy, interpret risk trade-offs, and make decisions that algorithms alone cannot resolve. The organization must treat agentic AI not as a replacement for human expertise but as a force multiplier, amplifying productivity while preserving control, trust, and alignment with enterprise values.

While many firms remain stalled at the personal productivity layer, Canopus demonstrates how shifting from speed at the edge to orchestration at the core begins with a unified workflow architecture—one that replaces fragmented administrative tasks with a governed, electronic trading environment capable of supporting a global hybrid workforce.

Edit Process: SLIP - KAVALA OIL - 915741DAA - Quotation Underwriting Task - Underwriter: Ramanna, Mahesh

James Fairgrieve
Tuesday, 13/04/2010

Settings Logout Help
Unreserve Close Process

Form

Instructors: Start Data Entry Task Contract Check Actions Request Aggregates Modelling Request Coverholder Approval

Profile Underwriting Info Notes Progress Tracking Referrals Contract Checks Subscribe Data Aggregates Coverholder Approval Peer Review Statistics History

Contact Info	
Placing Broker Contact Name	Mr John Smith
Placing Broker Email	john.smith@abcld.com
Policy Source	

Exposure	
Signing	20.00
Estimated Signed Line (%)	80.00%

Premium	
Premium Currency	
100% Estimated Gross Premium	1,500,000.00

Discounts	
Other Discount (%)	15.00%
Ultimate Total Discount (%)	5.00%

Additional Info	
Notice of Cancellation	30
Grading	Grade 2
<input type="checkbox"/> Reinsurance Required?	<input checked="" type="checkbox"/> Net Line
Syndicate Loss Ratio (ULR)	0.00
MOP Code	

Benchmarking	
Rating Model	
% of Benchmark, or % Loss Pick	0.00

Rate Changes	
Rate (%)	15.00%
Deductible (%)	22.00%
Commission (%)	5.00%
Limit (%)	25.00%
Underlying Rate (%)	22.00%
Other Rate (%)	13.75%

Contract Checks Comments		
Add Comment		
Comment	Added By	Comment Date
Data Please	dmSERVICE	08/04/2010 12:38:36


Documents
Related Processes

Business Process Management at Canopus

Canopus is one of the top 10 insurers in the Lloyd's insurance market with overseas operations in Bermuda, Singapore, Ireland, and Australia. Canopus has selected Enterprise Content Management (ECM) to improve processes, reduce filing costs, and support the implementation of electronic trading.

With technical advances in communications, the London insurance market has introduced an improvement initiative to reduce administrative overheads and improve servicing times. Electronic messages and documents are key elements of this improvement program, and participants in this market are tasked with adopting these new electronic solutions. Canopus was keen to support these improvement initiatives in readiness to trade with business partners on an electronic basis, while introducing more immediate efficiencies to its business processes.

Realizing the potential market demand for electronic trading, Canopus identified the need for a solution that would enable the organization to exchange information electronically with brokers and third parties around the world. They had already started work on a series of process models defining improved business functions and required a compatible workflow system to support these processes with minimal IT development effort. Adopting an ECM solution, Canopus relies on integrated business process and content management for long-term benefits such as improved visibility of 'work in progress', reduction in the effort needed for administrative tasks in the underwriting and claims areas, and general process improvements.



In prioritizing a platform that connects brokers and third parties through governed interfaces, Canopius has bypassed the brittleness of monolithic systems. Instead, they have created a modular, scalable foundation where human expertise is amplified by digital execution—proving that in the insurance market of the future, the winners will be those who can orchestrate complex global trades at the speed of thought.

Nested Orchestration as the Cognitive Operating Model

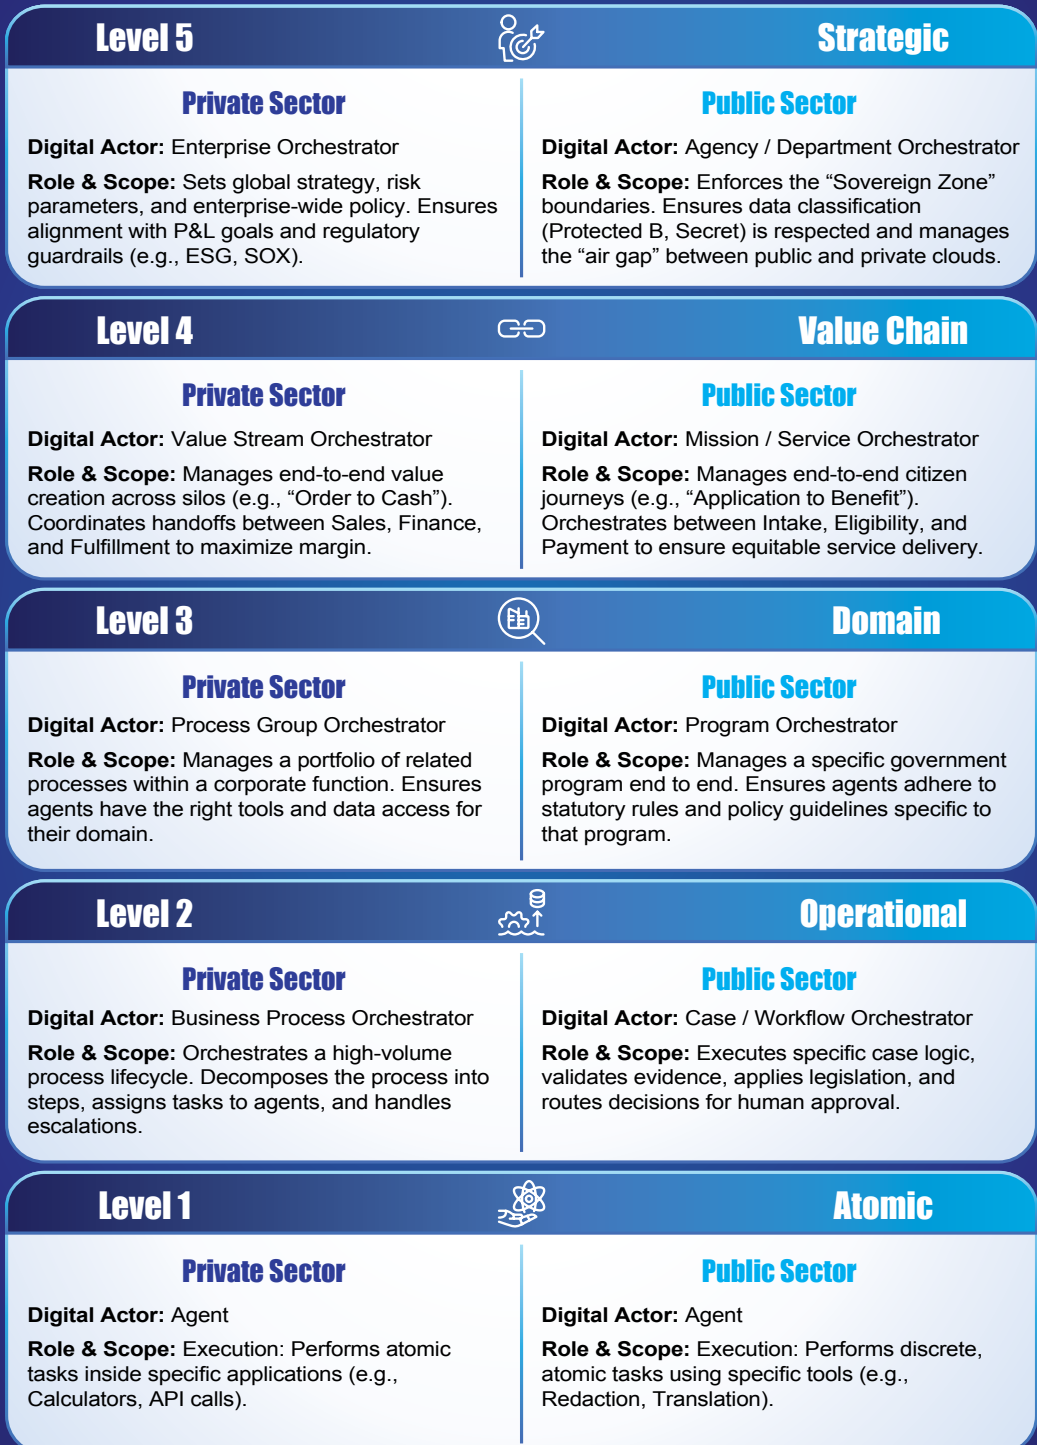
As we start to dive into the organizational structure to embrace human and digital workers, how the work is managed becomes a significant factor in ensuring your AI strategy's success. Orchestration of actions is critical, and more specifically, nested orchestration organizes intelligence into defined roles, each with bounded authority and measurable responsibility. Rather than relying on a single system to manage enterprise complexity, the organization distributes agency across orchestrated tiers. Each tier reduces complexity for the layer above it and enforces control for the layer below it.

This structure mirrors how mature enterprises manage scale, risk, and accountability. Analyst research reinforces this: among organizations identified as high-performance IT shops—those with strong alignment, deep trust, and high adaptivity—nearly 60% have placed their technology organization in charge of AI success. The operating model shift this chapter describes requires precisely this caliber of IT execution.

When these tiers of intelligence are mapped together—agents, orchestrators, policies, and workflows—they form the structural blueprint of the agentic enterprise. Later in this chapter, we describe this blueprint as the **Agentic Genome**.

From Org Charts to Orchestration

Redesigning the Operating Model for the Cognitive Enterprise



Level 1: The Atomic Unit – Anatomy of an Agent

To understand how this architecture works in practice, we begin at the smallest unit of intelligence: the individual agent. These are the specialized workers that execute discrete, atomic actions.

An agent is a highly specialized execution unit with narrow scope and high reliability. It does one thing and does it precisely. It doesn't know the broader context of the Order-to-Cash cycle. It only knows how to extract text from a PDF, calculate tax, reconcile transactions, validate identity, classify content, redact sensitive information, predict a risk score, or send an email.

It provides modular capability. Extract. Validate. Reconcile. Classify. Redact. Predict. Route.

Atomic or Level 1 agents don't manage workflows or own outcomes. They are reusable components of intelligence that are invoked by orchestrators with appropriate permissions.

For technology leadership, this modularity changes the economics of AI. Instead of building monolithic systems for each initiative, enterprises assemble reusable capabilities into orchestrated architectures. AI shifts from isolated projects to shared infrastructure.

Level 2: Operational – Process and Case Orchestrator (Operational Parent)

This is a level of execution intelligence above the functional layer of agents. It manages the execution of a single, specific workflow. Examples include loan origination, invoice processing, claims adjudication, contract review, court case management, or benefits determination.

As the "boots on the ground" commander, it's the Business Process Orchestrator in the private sector or the Case Orchestrator in the public sector.

This layer breaks down a complex goal into sequenced steps. It manages task ordering, decision branching, exception handling, retry logic, and human escalation. If an invoice is missing a purchase order number, it determines whether to reject the submission, request clarification, or escalate for review. If a benefits application lacks income verification, it initiates validation or routes the case appropriately.

With agentic AI, this is where operational performance improvements become measurable. Cycle times compress. Error rates decline. Service levels stabilize. Transformation becomes tangible rather than theoretical.

The transition from theoretical architecture to measurable ROI begins when atomic units of intelligence—like OCR and automated validation—are woven into a Level 2 operational orchestrator. In the case study below, Alabama Gas (Alagasco) illustrates this shift by transforming the fragmented Procure-to-Pay cycle into a unified, high-velocity engine that collapses hour-long manual tasks into mere minutes.

Alabama Gas

Alagasco

Alabama Gas Corporation (Alagasco) is the largest natural gas utility in the state of Alabama and is regulated by the Alabama Public Service Commission. With roots dating back more than 160 years, Alagasco today has operation divisions in Anniston, Birmingham, Gadsden, Montgomery, Opelika, Selma, and Tuscaloosa.

Alagasco is automating and optimizing its procure-to-pay process using an e-government solution. It integrates easily with the Company's ERP solution and effectively manages a large volume of invoices and purchase orders in its Procure-to-Pay process in accordance with its internal controls. The solution is used by all departments to make its procurement process more efficient and to help improve internal controls for compliance.

The solution allows Alagasco employees to work seamlessly between an e-government platform, Optical Character Recognition (OCR), Microsoft Office® Outlook®, and SAP. Using Vendor Invoice Management (VIM) software, users can code invoices and submit requestor approvals in a matter of minutes, a task that once took an hour or more to complete. Additionally, users have the ability to identify accrued liabilities; view all critical documents for a vendor electronically; and take advantage of special terms and discounting. Using the system, Alagasco has seen a reduction in the number of calls taken from vendors and improved its internal processes, allowing the company to increase productivity, reduce risk, and garner a better relationship with vendors.

Using Vendor Invoice Management (VIM) software, users can code invoices and submit requestor approvals in a matter of minutes, a task that once took an hour or more to complete.

By integrating Level 1 capabilities like OCR with a Level 2 VIM orchestrator, Alagasco has moved to a shared infrastructure of intelligence. This operational parent doesn't just manage the "boots on the ground" tasks; it enforces internal controls and auditability across the entire enterprise.

Level 3: Domain – Process Group and Program Orchestrator (Functional Parent)

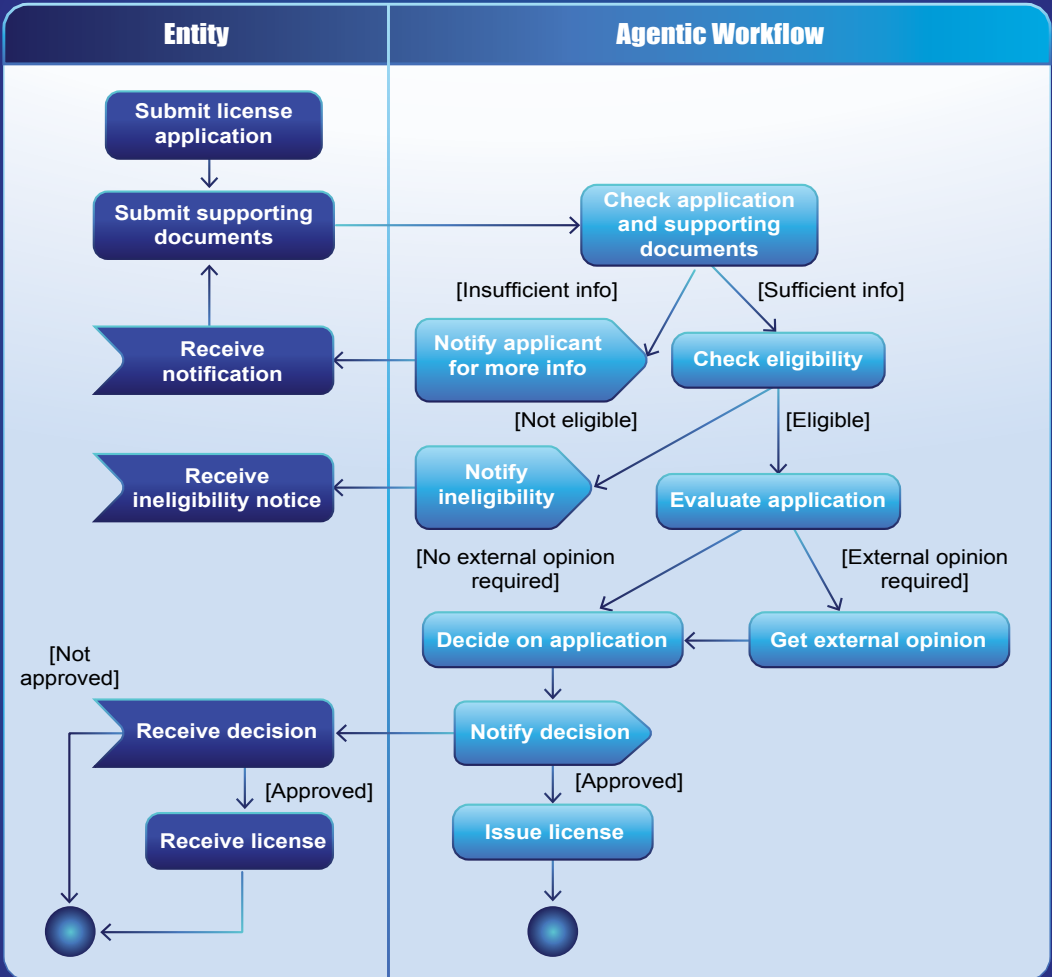
As guardians of logic, policy, and data, orchestrators at this level manage a portfolio of related capabilities (e.g., Accounts Payable, Social Services) and ensures domain-specific compliance. It acts as a firewall, embedding functional rules, regulatory constraints, and data permissions.

According to our levels of orchestration, Process Group Orchestrators in the private sector or Program Orchestrators in the public sector manage a specific functional domain such as Finance, HR, Legal, Supply Chain, Risk, Compliance, or Citizen Services.

This layer embeds domain specific logic, regulatory constraints, policy enforcement, and data permissions. It ensures that any intelligence operating within the domain adheres to internal controls and external regulations.

A marketing related capability cannot execute a general ledger transaction because domain boundaries are enforced structurally. A procurement process cannot override compliance thresholds without triggering policy logic. Access to sensitive records is controlled by design.

For the CIO and Chief AI Officer, this layer is where agentic AI becomes governable and auditable. For the Chief Data Officer, it protects data integrity, lineage, and contextual accuracy. Governance is no longer a guideline. It is encoded into execution.



This diagram illustrates an Orchestrated, Agentic License Renewal Workflow



The possibilities for employee self-service are almost endless. For example, if an employee sends a document attesting a change to their address or marital status, we aim to use AI-powered automation to update their employment record without a member of the HR team getting involved.

– Fausto Brembilla, SVP Head of Global HR Service Delivery

As a market leader in enterprise application software, SAP helps companies run better by redefining ERP and creating networks of intelligent enterprises that provide transparency, resiliency, and sustainability across supply chains. To deliver a high-quality HR experience to over 100,000 employees while streamlining General Data Protection Regulation (GDPR) compliance, SAP chose Content Management for its new SAP SuccessFactors Employee Central environment.

Previously, SAP relied on manual processes to ensure compliance with requirements such as the GDPR, sifting through thousands of HR documents to identify records for deletion. As it prepared to deploy SAP SuccessFactors Employee Central, the company saw an opportunity to automate these activities—strengthening compliance and boosting operational efficiency.

By deploying Content Management for SAP SuccessFactors, SAP has gained a modern, scalable document management platform to support its ongoing SAP SuccessFactors rollout—with powerful automation to help it accelerate compliance tasks. SAP is gaining the advantages of its next-generation Human Capital Management (HCM) solution while dramatically reducing the complexity of regulatory compliance. In the past, identifying the documents to dispose of under the GDPR was a laborious and time-consuming process that detracted from value-added activities. Today, they can define retention policies for each document type, delivering a highly automated approach to GDPR compliance across the entire SAP organization.

The SAP journey shows us the shift from labor scaling to intelligence scaling. By encoding GDPR retention policies into the architecture of their HR domain, SAP has ensured that as the company grows, its compliance burden does not grow with it. This is the hallmark of the Agentic Genome: moving away from manual searches toward a system where governance is an inherent property of the workflow.

Level 4: Value Chain – Value Stream and Mission Orchestrator (The Cross-Functional Parent)

If you consider that your organization delivers different outcomes via business processes, the concept of Parent Orchestrators really refers to the highest-level order of orchestration across the enterprise. This layer governs end-to-end outcomes that cut across functional silos. Examples include Order-to-Cash, Procure-to-Pay, Hire-to-Retire, Incident-to-Resolution, or Application-to-Benefit.

We refer to these Parent Orchestrators as Value Stream Orchestrators in the private sector or Mission Orchestrators in the public sector.

The Parent Orchestrator does not execute individual tasks. It owns the outcome. It maintains awareness of the entire lifecycle. It knows where a transaction resides, what dependencies remain unresolved, and what risks threaten completion. It coordinates across domains to ensure continuity of flow.

For executive leadership, this layer operationalizes strategy. It provides visibility into systemic bottlenecks and ensures that cross-functional performance aligns with enterprise objectives.

Level 5: Strategic – Enterprise / Agency Orchestrator (The Strategic Parent)

Think of the Enterprise Orchestrator as the conductor of a vast, complex symphony, the CEO, or the Board of the digital workforce. Its primary function is not just to set high-level policy (which is often codified into lower-level agents) but to ensure alignment and orchestration health.

While governance rules (like GDPR or spend limits) are often built directly into the individual agents and lower-level orchestrators for real-time enforcement, the Enterprise Orchestrator manages the relationships and conflicts between the major Value Streams.

It answers the question: *Are our disparate business units working together or pulling apart?* This layer prevents “local optimization,” where one department wins at the expense of the whole, by undermining enterprise objectives.

From Labor Scaling to Intelligence Scaling

The economic implications of nested orchestration are significant. Traditional enterprises scale by adding labor. Growth often implies proportional increases in headcount and cost. An agentic enterprise scales by adding intelligence.

Once orchestrated layers are established, incremental volume does not require proportional human expansion. Decisions can be made continuously rather than sequentially. Policy enforcement becomes structural rather than procedural. Failures can be isolated within specific layers and resolved without destabilizing the entire system.

Human expertise does not disappear. It shifts upward toward oversight, judgment, and strategic direction.

For executive leadership, the strategic question becomes whether the organization can increase throughput, resilience, and compliance posture without increasing operational volatility. Nested orchestration provides a pathway to that outcome.

Humans, Agents, and the New Structure of Work

The agentic enterprise will not be defined by AI replacing workers one-for-one. As explained in this chapter, work itself is reorganized. Agentic systems decompose complex roles into smaller executable components, allowing work to be performed in parallel, continuously, and at dramatically larger scale.

In operational terms, a single well-designed agent can absorb tasks previously distributed across five to twenty-five people performing partial roles. The multiplier appears high, but the economics are straightforward. Agents operate continuously, twenty-four hours a day, with near-zero marginal cost, executing tasks in parallel across thousands of transactions, without fatigue and with consistent accuracy.

This does not eliminate human roles. Rather, it shifts where humans operate in the system.

Humans increasingly decide, supervise, resolve ambiguity, and maintain accountability, while agents execute, analyze, monitor, simulate, and coordinate. As agent populations grow, the human role concentrates around judgment, accountability, and governance (see the section below).

The result is a structural change in how organizations scale. In traditional enterprises, capacity scaled primarily through hiring. Organizations expanded by adding more people to perform more work. In the agentic enterprise, this relationship changes.

The human workforce becomes a slow-moving stock variable, while the agent population becomes a fast-scaling flow variable. Hiring cycles typically operate on the scale of months or years. Recruiting, training, and organizational integration take time. Agent deployment cycles operate on the scale of hours or days. New agents can be instantiated, configured, and integrated into workflows almost immediately.

This difference fundamentally alters competitive dynamics. Organizations that can scale digital execution rapidly gain structural speed advantages over those constrained by traditional workforce growth.


The New Constraint on Work

In mature deployments, organizations may operate tens of thousands of agents coordinated by a few thousand orchestrators alongside a largely stable human workforce. Work is no longer primarily constrained by the number of available people.

Instead, it becomes constrained by the infrastructure that enables digital execution:

- Data readiness
- Compute budget
- Orchestration design
- Governance, risk, and sustainable trust

These elements determine how quickly an organization can expand its agent population and how safely that expansion can occur. Organizations that solve these constraints early—establishing governed data environments, scalable orchestration platforms, and trusted control frameworks—gain a lasting structural advantage. They become not just more efficient, but structurally faster competitors, capable of scaling execution at a pace traditional organizational models cannot match. This is illustrated in the following image.



Agent vs Human Scaling Curve (2025–2045)

In the agentic enterprise, work no longer scales primarily through headcount. It scales through orchestration—coordinating humans, agents, and information into a continuously operating system of execution.

Governance, Risk, and Sustainable Trust

Agentic AI is ultimately a governance architecture decision.

Unstructured deployments introduce compliance exposure, data leakage risk, decision opacity, and unpredictable behavior. Intelligence that operates without bounded authority cannot be trusted at scale.

Nested orchestration addresses this structurally. Authority is scoped at each layer. Decisions are logged in context. Permissions are enforced by design. Escalation paths are defined. Explainability becomes embedded in workflow rather than retrofitted after the fact.

Trust becomes an architectural property of the system.

Critically, enterprise data consistently shows that organizational readiness scores the lowest of all AI success factors—typically 50–60%—while technical feasibility and market timing score 80–95%. The gap is not capability; it is coordination. Trust at scale requires cross-functional alignment across technology, business, compliance, and security leadership—not just technical governance controls. For regulated industries and public institutions, this distinction is decisive.

AI that cannot be governed cannot be scaled. AI that cannot be audited cannot be defended. Nested orchestration transforms AI from a potential liability into a controllable asset. It enables demonstrable compliance, defensible decision trails, and structured accountability across the enterprise.

From Functions to Flows

Perhaps the most profound transformation is organizational rather than technical.

Forrester’s “State of AI Survey” confirms the structural problem: 53% of enterprises place AI leadership within the technology organization, while only 7% assign it to business functions. CIOs and CTOs hold 44% of AI technology strategy responsibility but only 29% of AI business strategy.⁷ This imbalance explains why AI initiatives succeed technically but fail to transform operations—the business side is not at the table. Org charts still exist in the agentic enterprise, but they no longer define how work truly happens. The real operating model resides in orchestrated layers of intelligence.

This shift from functions to flows also means risk is monitored across entire lifecycles rather than within isolated departments, allowing earlier detection and more coherent intervention.

The Emergence of the Agentic Genome

When thousands of agents operate within nested orchestration that is governed, bounded, and aligned to outcomes, the enterprise evolves into something fundamentally new. No longer defined primarily by reporting lines or application stacks, it becomes a coordinated network of intelligence capable of sensing, reasoning, deciding, and acting across its own value creation mechanisms.

Agents execute atomic actions.

Operational orchestrators manage workflows.

Domain orchestrators enforce policy and compliance.

Value stream orchestrators coordinate outcomes across functions.

Enterprise orchestration aligns the system with strategy.

Together, these layers form the structural blueprint of the agentic enterprise. We refer to this blueprint as the **Agentic Genome**.

Just as biological genomes encode the instructions that govern how living organisms function, the Agentic Genome encodes how intelligence operates inside the enterprise—how agents are structured, how decisions flow, how governance is enforced, and how outcomes are achieved.

The enterprise is no longer defined primarily by its org chart or application landscape. Instead, it becomes a coordinated network of intelligence capable of sensing, reasoning, deciding, and acting across its own value creation mechanisms.

In this model, intelligence becomes infrastructure.

The shift from org charts to orchestration is not cosmetic—it is foundational. The organizations that recognize this will not simply deploy AI. They will redesign themselves around it.

Executive Implications

CAIO

Agentic transformation requires coordinated governance. The CAIO bridges technology, business, compliance, and security leadership to ensure AI moves from technical pilots to an enterprise operating model.

CIO

The CIO's role shifts from executing AI projects to enabling an enterprise agentic platform—providing the standards, architecture, and security foundations that allow business units to safely deploy intelligent workflows.

CFO

Incremental productivity gains rarely justify escalating AI spend. Structural operating model redesign is required to produce measurable ROI and margin expansion.

CHRO

AI cannot remain a personal assistant experiment. The workforce must evolve toward human-agent collaboration, with literacy and confidence as core retention drivers.

CDO

Faster outputs amplify underlying data quality issues. Sustainable performance requires governance and accuracy embedded into execution layers.

COO

True operational improvement does not come from faster tasks, but from redesigning how work flows across functions. Cycle time reduction depends on systemic coherence.

In the next part of the book, we introduce the “Agentic Genome Map” as an AI agent model and examine enterprise functions according to this model, orchestrated across organizational roles, departments, and industry domains.

Part 2

Enterprise Functions and Agents Model

02

Chapter Five

What Is an Agentic Genome Map? (and Why You Need One)

If Chapter 4 defined the operating model of the agentic enterprise, this chapter introduces the Agentic Genome Map—a visual and conceptual operating model for how intelligence is structured and executed inside an enterprise. It explains how agents, orchestration, and enterprise information systems work together to produce coordinated, governable AI at scale.

The Agentic AI Genome: DNA of the Cognitive Enterprise

The **Agentic AI Genome** is the organizing operating model for the agentic enterprise. It defines how autonomous agents, orchestration layers, governed enterprise information, and human-in-command oversight are structured into a cohesive system that can sense, decide, and act across workflows at scale.

Rather than describing a single application or model, the Agentic AI Genome represents a repeatable architectural pattern. At its foundation lies trusted enterprise content and systems of record—the institutional memory that grounds agent reasoning. Above this sits the agent layer, where domain-specific agents perform specialized tasks. The orchestration layer coordinates multi-step workflows, enforces policy, manages escalation paths, and ensures execution integrity. At the outer layer, human command defines objectives, autonomy boundaries, risk thresholds, and ethical guardrails.

The Genome is “genetic” in that it is replicable across domains—finance, HR, supply chain, operations—while maintaining consistent governance, lifecycle controls, and sovereign deployment practices. It enables organizations to scale intelligent execution without fragmenting architecture or diluting accountability.

In practical terms, the Agentic AI Genome integrates:

- Governed content and enterprise systems as the agentic substrate
- Specialized agents that reason and act within defined autonomy tiers
- Orchestration and control planes that enforce policy and auditability
- Continuous observability and lifecycle management
- Human-in-command leadership to define intent and retain accountability

The result is not isolated AI capability, but an enterprise-wide execution fabric where speed and sovereignty evolve together.

Mapping the Invisible System

We understand that success lies in nested orchestration, a hierarchical system where strategy flows down and execution flows up. But a practical executive question remains: How do you actually visualize, design, and manage this complexity? How do you ensure that the thousands of agents operating within your finance, HR, and supply chain functions are not just active, but aligned? How do you audit a digital workforce that makes decisions in milliseconds?

You need a map.

When organizations digitized processes, they created visibility into workflows. When organizations deploy agentic AI, they introduce something fundamentally different: an invisible, distributed layer of decision-making and execution. Agents consume data for context, operate across systems, trigger actions, and influence outcomes. Without a unifying model, this intelligence layer quickly becomes opaque. And opacity is the enemy of scale.

Just as the Human Genome Project mapped the building blocks of biological life to understand health and disease, the Agentic Genome Map represents the building blocks of digital intelligence to understand enterprise performance, risk, and capacity.

More than serving as a technical diagram, the map is a strategic blueprint. It enables the CIO, CDO, and CEO to see, often for the first time, the direct relationship between a high-level business objective and the specific digital actors executing to achieve it.

The Agentic Genome Map is more than an architectural diagram; it is the strategic blueprint that allows a modern institution to visualize how invisible threads of data become concrete financial actions. For this ambitious Saudi Arabian bank, building this digital backbone was the essential first step in mapping its path from manual, batch-processed silos to a real-time execution fabric where compliance and customer service operate as a single, coordinated organism.



A Saudi Arabian Bank



As a digitally ambitious financial institution, this Saudi Arabian bank operates in an environment defined by rising customer expectations, stringent regulatory requirements, and growing transaction volumes across mobile and online channels. By digitizing key workflows and integrating content, process, and policy systems on a unified digital backbone, the bank has laid the foundation for intelligent automation that closely mirrors agentic patterns—where data flows uninterrupted through governed pipelines and decision logic can be embedded directly into execution.

This digital strategy aligns with embedding intelligent capabilities into core processes: end-to-end workflows that once relied on manual checks, paper-based verification, and periodic batch processing are now primed for real-time, context-aware orchestration. In this context, agents could operate on continuously ingested data to interpret customer information, assess risk against policy constraints, and trigger appropriate actions—such as initiating compliance checks, updating customer records, or escalating high-risk signals to human reviewers.

Key operational improvements include end-to-end workflow automation across onboarding, KYC, and transaction monitoring; real-time compliance and risk checks embedded into execution; continuous orchestration across channels (mobile, web, branch), and structured handoffs to human reviewers with full audit trails and explainability. The impact is a more resilient, responsive operating model: faster customer experiences, stronger compliance posture, and scalable operations that improve speed, accuracy, and trust across the banking value chain.

By treating its digital infrastructure as a repeatable, “genetic” pattern, the institution has created a model where speed and sovereignty evolve together. This is the Agentic Genome in action: a more resilient, responsive operating model that doesn’t just process transactions faster, but creates a defensible environment where every millisecond of decision-making is aligned with the bank’s ultimate mission of trust and stability.

Defining the Agentic Genome Map

The Agentic Genome Map is a multidimensional visualization of your organization's total operational reality. It models the intersection of:

- **Vertical Context:** Industry dynamics, regulations, and market structure
- **Horizontal Functions:** Finance, HR, supply chain, legal, and IT
- **Execution Intelligence:** Orchestrators, processes, and agents
- **Governance and Risk:** Controls, permissions, and oversight
- **Value Flows:** Cost, revenue, resilience, and capacity

However, viewing this map solely as a picture of “intelligence” is insufficient. It is a dynamic visualization of execution, effectiveness, and governance. It maps the metabolic rate of your business, showing where data (fuel) is being converted into decisions and actions (motion) by digital workers (agents).

This map serves three critical functions for leadership:

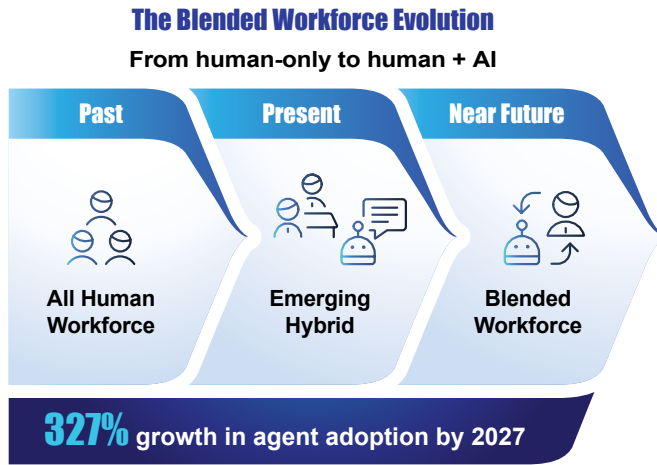
1. **Visibility:** It reveals where AI is deployed, what data it touches, and who owns it. What was previously fragmented becomes observable.
2. **Governance:** It explicitly links agents to compliance regulations and human oversight. Governance shifts from reactive enforcement to architectural design.
3. **Value:** It connects operational activity to realized ROI, preventing the “pilot purgatory” where AI projects stall. Every agent exists for a reason, and every capability must produce value.

The Digital Workforce: Agents as Digital Actors

AI agent adoption is expected to increase 327% over the next two years; digital labor is the future and its integration into organizational roles will be critical for success.⁸

AI agents will soon be embedded across every business function and industry, reshaping how work is organized and executed. For managers, this means leadership will no longer be limited to guiding people alone but will extend to overseeing a blended workforce of humans and AI agents working side by side.

In the agentic enterprise, we must stop thinking of agents and orchestrators merely as code or software scripts and start treating them as Digital Actors or Digital Workers.



Digital and Human Workers Collaborate in a Blended Workforce⁹

When you introduce an agent into the enterprise, you are effectively “hiring” a digital employee. This analogy is not just rhetorical; it provides the robust management framework required for scale. Just like human employees, digital actors require the full lifecycle of HR management, adapted for silicon rather than biology.

1. Role Definition and Job Descriptions

Just as you would never hire a human without a job description, you should never deploy an agent without one. Ambiguity at scale becomes risk. The Genome Map requires a “Digital Job Description” for every agent, defining its purpose, scope of authority, required capabilities (Tools), reporting line (Orchestrator), and success metrics.

2. Identity and Access Management (IAM)

Digital actors need digital identities to access critical information, systems, and processes. Permissions define power and boundaries define safety. We must apply the principle of least-privilege access to agents just as we do to humans. An agent hired to schedule meetings should not have permissions to read payroll databases. The Genome Map becomes a structural enforcement mechanism for these identity boundaries.

3. Training and Context Alignment

Human workers undergo onboarding; agents undergo fine-tuning and context engineering. Humans need upskilling; agents need model updates and retraining on new data. As business conditions evolve—new policies, new data structures, new regulations—agents must adapt. The Genome Map tracks the “training status” of your digital workforce to ensure their skills remain relevant.

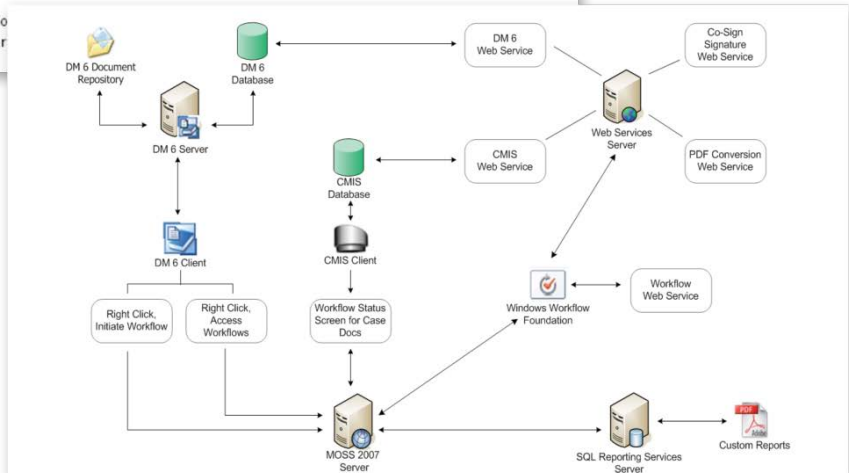
4. Performance Management

We measure human performance against KPIs. We must do the same for Digital Actors. Key questions include accuracy, latency, drift, exception rates, and value contribution. The Genome Map integrates Performance Management and Measurement frameworks to ensure the digital workforce is contributing to the P&L, not just consuming compute cycles.

If introducing an agent into the enterprise is akin to hiring a digital employee, then the first step in the Agentic Genome Map is identifying where manual friction creates the most significant drag on human talent. In the following case study, the European Court of Human Rights illustrates this transition by redefining its internal approval processes as a series of automated workflows that treat case management not as a clerical task, but as a coordinated digital role.

European Court of Human Rights

The European Court of Human Rights (ECHR) is part of the Council of Europe, an international intergovernmental organization that was established in 1949. Currently, the Council is made up of 47 member states that have all signed up to the system of human rights protection under the European Convention on Human Rights.



Streamlining Process at ECHR

Over the last few years the Court's caseload has exploded, with the number of applications to the Court growing from 14,000 to over 54,000. Recognizing that it was time to streamline its internal approval processes, the Court's IT Department developed an in-house automated workflow solution. The workflow helps manage the approval process for committee and chamber cases and provides the Court with a mechanism that streamlines its case management processes, further enhancing the productivity of the legal divisions.

Overall user sentiment is that the system saves time, is easy to use, has led to less work for the legal assistants, and is more streamlined than sending the committee notes by paper. Division assistants highlight the fact that the solution automatically tracks the route of the workflow, and that the dynamic reports make it easy for the divisions and sections to find out what stage a workflow is at. With the workflow solution for committee notes successfully underway, the Court plans to begin piloting its workflow for chamber notes.

ECHR has moved beyond paper-based limitations to create the observability required for a blended workforce. As the Court implements workflows for chamber notes, it isn't just installing software; it is mapping the metabolic rate of its legal divisions for a future where humans and agents uphold human rights side-by-side.

The Orchestration Hierarchy: Governance by Design

The vertical spine of the Genome Map is the Orchestration Hierarchy. As discussed in Chapter 4, "From Org Charts to Orchestration," this is not a flat network. It is a parent-child relationship structure that ensures security, modularity, control, and a clear scope of authority. For review, the (simplified) structure is:

Level 1: Agent

Level 2: Business Process Orchestrator

Level 3: Process Group Orchestrator

Level 4: Value Stream Orchestrator

Level 5: Enterprise Orchestrator

This structure enhances security. It creates "blast radius" containment. If a Business Process Orchestrator in the shipping department fails, it doesn't bring down the entire Supply Chain Value Stream. The parent orchestrator simply detects the failure and routes around it or alerts a human.

Furthermore, the structure allows for modular upgrades. You can swap out a Level 1: Atomic Agent (e.g., a translation tool) for a better one without rewriting the logic of the entire enterprise. The Business Process Orchestrator simply calls the new agent, and the work continues uninterrupted.

Let's take a more granular look at the foundational level—the Agent—so we can delve into how the Genome Map works in practice.

Level 1: Agent

In our map, an agent is defined by far more than its code. To be “enterprise ready,” every agent on the map must be characterized by seven critical attributes. These attributes turn a “bot” into a managed asset:

1. Context Sources

An agent is only as good as what it knows. The Genome Map explicitly lists the content sources required to create, train, operate, and govern the agent. Does it need access to the CRM? The ERP? A static PDF library of policy documents? Defining these dependencies makes integration intentional and ensures IT knows exactly what “pipes” need to be connected.

2. Tool Access

Intelligence without the ability to act is merely analytics. To complete its mission, an agent must be granted specific operational capabilities, or “Tools,” to interact with the enterprise environment. The Genome Map explicitly catalogs the functional tools an agent is authorized to use—API connectors, interpreters, browsers, and critically, governed Model Context Protocol (MCP) integrations. MCP is emerging as the standard mechanism for connecting agents to enterprise content. However, a critical design principle applies: the value must live in the governance layer, not in the connection itself.

If MCP servers are simply pipes that pass content, they are easily commoditized. If they are governed content delivery services—with access controls enforced, compliance metadata attached, audit trails generated, and sovereignty routing built in—they become defensible infrastructure that no competitor can easily replicate. The Genome Map must therefore catalog not just which tools an agent can use, but what governance controls are embedded in each connection. Defining Tool Access ensures that IT and security teams can strictly enforce the principle of least privilege, guaranteeing, for example, that an agent designed to extract invoice data cannot accidentally authorize a payment.

3. Data Types

Not all data is equal. The map categorizes the fuel the agent consumes as human generated (high context, high nuance), machine generated (high volume, high structure), and transactional and business network data (high precision, zero tolerance for error). Understanding data types is crucial for selecting the right underlying model (e.g., an LLM for human text vs. a regression model for machine logs) and determining risk posture.

4. The Human Oversight Owner

Autonomy without accountability is instability. Every agent must have a named Human Oversight Owner operating within a framework architected by the Chief AI Officer (CAIO). The CAIO has emerged as the primary architect of AI governance, working in concert with the CIO and CISO to define the guardrails that individual owners must enforce.

The Oversight Owner has the oversight and responsibility to ensure each Agent is monitored for:

- **Calibration:** Providing feedback to improve the agent's model
- **Exception Management:** Handling the "edge cases" the agent cannot resolve
- **Ethical Boundary Enforcement:** Ensuring the agent does not drift into biased or unsafe behavior

There is no autonomy without accountability.

While every agent requires a single named owner, the governance framework within which these owners operate is inherently multi-stakeholder. Current data shows that AI governance depends on coordinated leadership across senior roles: the Chief AI Officer (28%), CIO (24%), CISO (22%), CTO (20%), Chief Data Officer (17%), with business stakeholders and developers each contributing 13%.¹⁰ This is not a committee for each agent—it is a shared governance architecture that individual owners enforce. The CAIO defines the guardrails; the CIO provides the platform; the CISO secures it; and the CDO ensures data integrity. The Genome Map must reflect this: each agent's oversight owner operates within a framework that only works when these roles are coordinated.

5. Related Regulations

Agents inherit compliance obligations from the data and decisions they influence. The map links every agent to the compliance frameworks that govern its data. An agent handling employee health records is tagged with "HIPAA" or "GDPR." An agent handling defense supply chains is tagged with "ITAR." This allows the Governance and Compliance Agents (at the Process Group level) to automatically audit compliance based on the agent's tags.

6. Risk Assessment

Not all agents carry equal impact. Every agent must be classified according to a risk framework aligned with the EU AI Act—whose high-risk rules take effect August 2026 with fines up to 7% of global revenue.¹¹ The Genome Map adopts the Act's four-tier classification:

- **Unacceptable Risk:** Agents that manipulate human behavior, exploit vulnerabilities, or perform real-time biometric identification in prohibited contexts. These are banned outright.
- **High Risk:** Agents that touch PII, make financial decisions, control physical machinery, or operate in domains such as healthcare, law enforcement, or critical infrastructure. These require conformity assessment, risk management systems, human oversight, and complete record keeping.
- **Limited Risk:** Agents that handle internal operational data, proprietary project timelines, or internal communications. These require transparency obligations and periodic audit.
- **Minimal Risk:** Agents that summarize public news feeds or schedule internal meetings. These may operate with greater autonomy under standard monitoring.

This regulatory alignment is not aspirational—it is immediately actionable. For organizations subject to the EU AI Act, NIST AI RMF, or ISO 42001, every agent on the Genome Map must carry a risk classification that drives its governance controls.

7. Latent Value and ROI

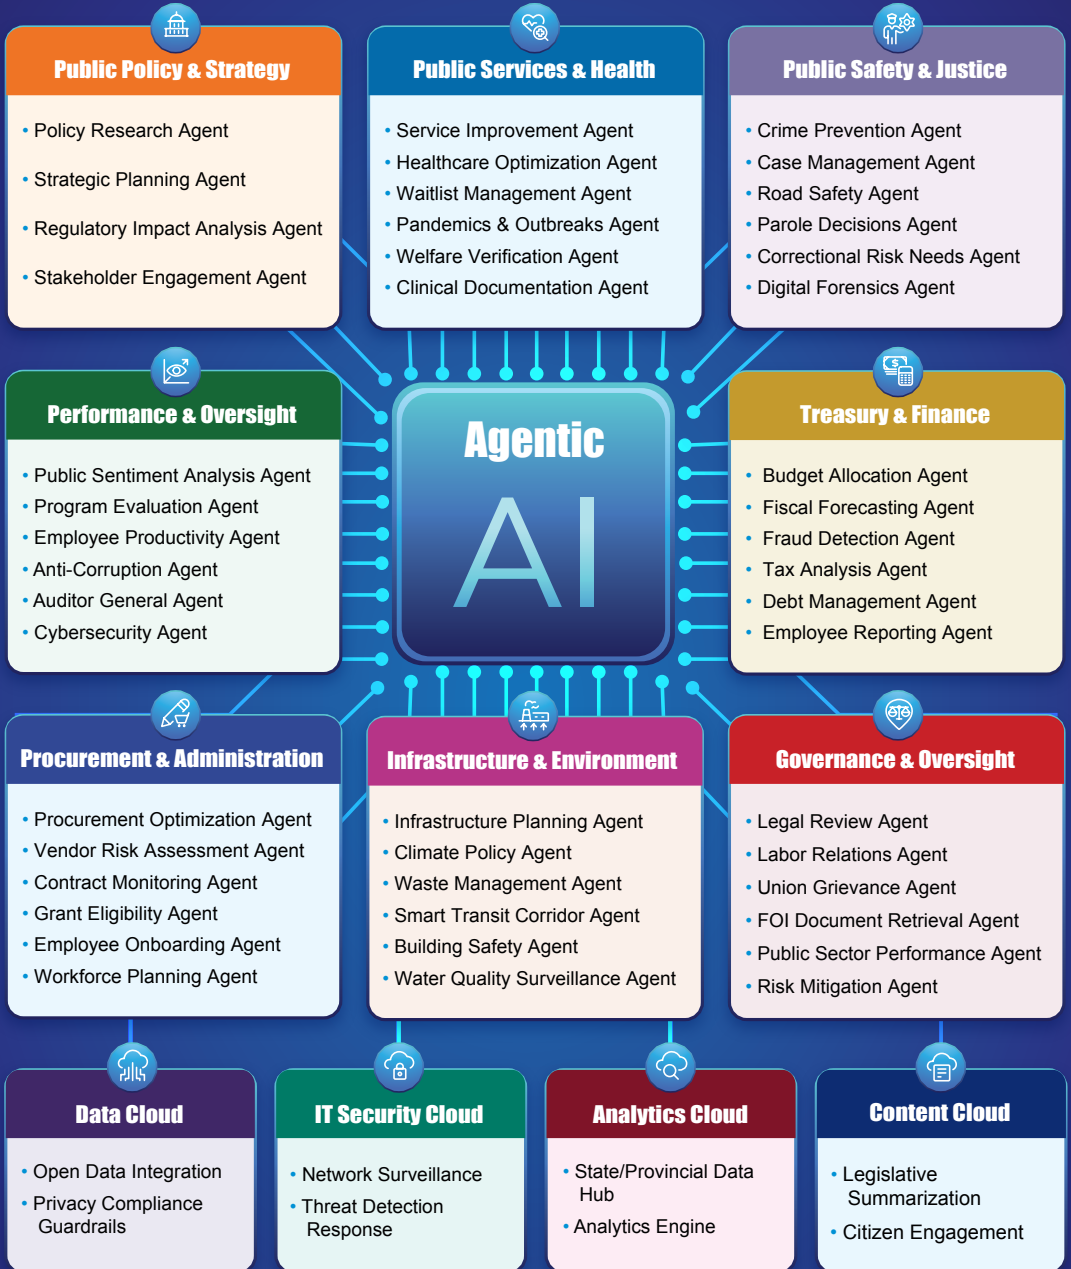
Finally, the Map quantifies the why. Every agent must answer a fundamental economic question: Why does this agent exist? Is it cost reduction (hours saved)? Is it revenue generation (faster speed to lead)? Is it risk or cost avoidance (fewer compliance fines)? By mapping this at the atomic level, the CFO can roll up the ROI of the entire AI portfolio with precision.

The Agentic Genome Map – The Enterprise's Digital Nervous System

The Agentic Genome Map is a structural model for how intelligence operates, scales, and remains governable within the cognitive enterprise.

Like a biological nervous system, the architecture distributes cognition across specialized layers. Strategic Orchestrators define intent and constraints. Value Chain and Domain Orchestrators translate objectives into coordinated operational direction. Operational Orchestrators manage execution dynamics. Atomic Agents perform precise actions.

State / Provincial Government Agentic Genome Map



A Genome Map Featuring a Sampling of Agents (Without Requisite Layers of Orchestration)

Each layer is distinct. Each role is bounded. Each decision remains contextually aligned.

The model ensures that enterprise intent does not fragment as it propagates through systems, processes, and workflows. Governance remains intact. Coordination remains fluid. Execution remains adaptive.

This is not a visualization of AI components. It is a blueprint for how an enterprise thinks and acts as a unified system.

Putting It Together: The Chain of Execution

To understand how the Genome Map works in practice, let us walk through the concrete example of invoice processing. We'll look at how a single Agent propagates value up the chain through various orchestrators.

The Scenario: Invoice Processing

Level 1: The Agent (The Digital Doer)

We deploy an Extraction Agent. Its context source is the vendor email inbox, with a data type that is an unstructured PDF and a Human Oversight Owner, the AP Manager. The Agent's value is eliminating manual data entry.

Role and Scope: The Agent reads a PDF invoice, extracts the Vendor Name, Amount, and Date, and structures it into a JSON (JavaScript Object Notation) file.

Level 2: Business Process Orchestrator (The Digital Manager)

The Extraction Agent reports to the Invoice Processing Orchestrator.

Role and Scope: This orchestrator receives the structured data. It contacts a Validation Agent to check that the math adds up and a Matching Agent to check the ERP for a corresponding Purchase Order. If everything matches, it approves the invoice. If not, it routes the exception to the Human Oversight Owner.

Level 3: Process Group Orchestrator (The Digital Director)

The Accounts Payable Orchestrator oversees the flow of thousands of invoices.

Role and Scope: This orchestrator monitors liquidity needs and ensures that the Invoice Processing Orchestrator is prioritizing vendors with early payment discounts. It manages the relationship between the Invoice workflow and the Vendor Onboarding workflow.

Level 4: Value Stream Orchestrator (The Digital VP)

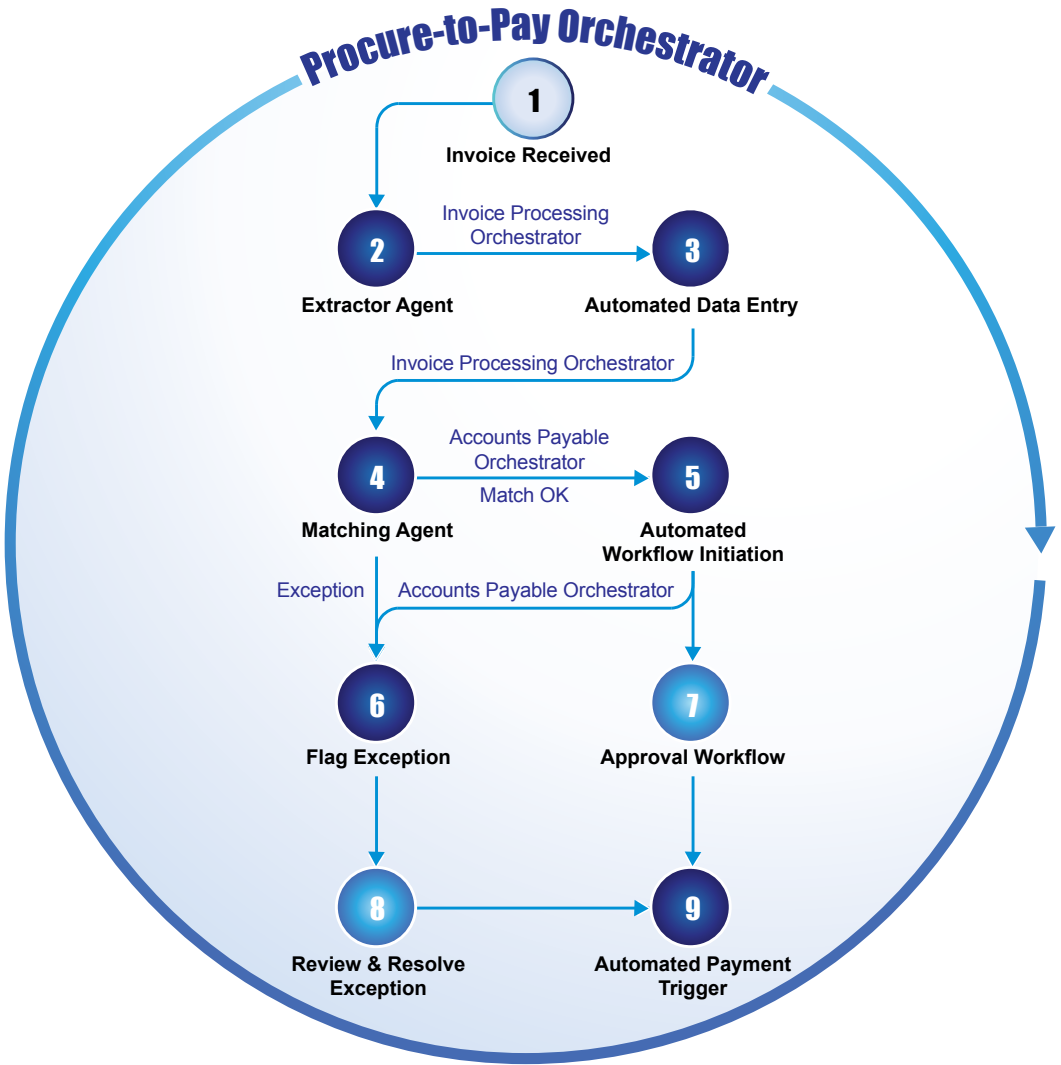
The Procure-to-Pay Orchestrator governs the entire cycle.

Role and Scope: This orchestrator ensures what the company is purchasing (Procurement) aligns with what it is paying for (Accounts Payable) and that the goods are actually arriving (Inventory). It optimizes working capital across the entire supply chain.

Level 5: Enterprise Orchestrator (The Digital C-Suite)

The Governance Orchestrator oversees it all.

Role and Scope: While the Procure-to-Pay Value Stream is efficiently optimizing working capital, the Enterprise Orchestrator ensures this isn't happening in a vacuum. It monitors the health and alignment between this stream and others, such as Order-to-Cash or Treasury Management. If the Procure-to-Pay stream is aggressively paying invoices to capture early discounts (draining cash), but the Treasury stream is signaling a liquidity crunch due to slow collections, the Enterprise Orchestrator detects this systemic friction. It intervenes not by stopping the invoice, but by re-calibrating the objectives for the Value Stream Orchestrators, perhaps temporarily deprioritizing early payments to preserve cash flow until collections stabilize. It ensures the efficiency of one department doesn't accidentally harm the health of the whole.



Agentic Invoice Workflow

As KRAMSKI manufactures components to a tolerance of a thousandth of a millimeter, its transition to a more autonomous, intelligent enterprise begins with the high-precision deployment of Level 1 extraction agents and Level 2 process orchestrators to eliminate the manual friction of paper-based approvals.

KRAMSKI



With intelligent automation from Invoice Management for ERP Solutions, we can meet the business requirements while keeping our back-office lean.

— Andreas Heyde, Senior Project Manager, KRAMSKI

For almost half a century, KRAMSKI has been producing precision-engineered components for a wide range of sectors, including the automotive, electronics, and medical industries. The company manufactures its products to extremely fine tolerances—in many cases, down to just a thousandth of a millimeter.

As the company continued to grow, it aimed to keep its back-office lean and efficient while maintaining the highest quality standards. However, manual, paper-based processes presented a challenge. In the past, KRAMSKI manually circulated paper invoices for signed approval before accounts payable (AP) teams manually entered data from paper and electronic invoices into their ERP for posting and payment—a time-consuming and labor-intensive process. To maintain the quality and responsiveness of its back-office services, KRAMSKI began looking for a way to streamline these workflows.

KRAMSKI searched for a single vendor to deliver all its automation goals: I2P, P2P, O2C, and e-invoicing. Using intelligent automation and Vendor Invoice Management, KRAMSKI now automatically extracts data from paper and electronic invoices and integrates it seamlessly into its ERP system. Automated workflows then guide each document through review and approval stages. Seamless integration has let the company streamline document-related transactions, manage operational costs, and lay the foundation for scalable business growth as the solution is gradually rolled out.

By automating the atomic task of extraction, KRAMSKI has enabled a higher-level operational parent to manage exceptions and ensure domain-level compliance. This transition ensures that the company's administrative "metabolic rate" now matches the speed and precision of its world-class manufacturing lines.

Public vs. Private: Same Genome, Different Expression

The beauty of the Agentic Genome Map is its universality. The structural hierarchy applies equally to a Fortune 500 retailer and a Federal Agency. The structure of intelligence remains constant; only the terminology and the value drivers change, as exemplified below.

In the private sector:

- **Actors:** Value Stream Orchestrators managing Order-to-Cash
- **Value:** Measured in margin, revenue growth, and shareholder value
- **Focus:** Competitive advantage and efficiency

In the public sector:

- **Actors:** Mission or Service Orchestrators managing Application-to-Benefit
- **Value:** Measured in citizen trust, service equity, and mission fulfillment
- **Focus:** Security, stewardship, and compliance

In both worlds, the Genome Map functions as the essential “digital twin” of enterprise intelligence. It transforms the abstract chaos of AI adoption into a structured, governable, and measurable science. Agentic Genome mapping transforms AI adoption from experimentation into engineering, from activity into accountability, and from promise into performance.

Without a Genome Map, agentic AI introduces complexity faster than organizations can govern it. With a Genome Map, intelligence becomes visible, governable, measurable, and optimizable. It becomes infrastructure. It becomes strategy. It becomes the foundation of the agentic enterprise.

An Agentic Genome Map is not optional documentation. It is the management system for intelligence at scale.

Executive Implications

CAIO

The CAIO defines the governance framework that makes the Genome Map operational—establishing risk classification, oversight standards, and cross-functional coordination that connects AI architecture to business accountability.

CIO

The Agentic Genome Map becomes the operating blueprint for scaling AI safely, making distributed intelligence visible, governable, and modular across the enterprise architecture.

CFO

Mapping every agent to value streams and ROI transforms AI from exploratory spend into a managed portfolio of digital assets with measurable economic impact.

CHRO

As agents become digital workers, leadership, accountability, and performance management must extend beyond people to include the governance of a blended human–AI workforce.

CDO

Without a genome map linking agents to data sources, lineage, and regulatory constraints, enterprise AI quickly becomes opaque and non-defensible under audit.

COO

Operational coherence depends on the orchestration hierarchy, ensuring thousands of agents execute locally while remaining aligned to enterprise-level objectives and constraints.

In the next chapter, we bring the Agentic Genome Map to life by showing how agents and orchestrators operate across core enterprise functions—from the boardroom to finance, HR, sales, IT, operations, and legal. You'll see how the same genome expresses itself differently across domains, translating the operating model into practical, high-impact use cases.



Chapter Six

Applying Agentic Capabilities to the Enterprise

In practice, agentic AI manifests not as a monolithic system, but as domain-specific agentic applications embedded across executive, financial, operational, and regulatory functions. In this chapter, we illustrate how agentic capabilities express themselves across departments—each with distinct goals, workflows, and governance requirements, yet all sharing the same underlying genome: agents coordinated by orchestrators, grounded in enterprise information, and governed by humans-in-command.

In the agentic enterprise, point solutions may scale technically but without active business leadership, intelligence does not translate into value. Teams deploy isolated agents, pilots proliferate, and models improve—yet the enterprise lacks a coherent way to organize, govern, and scale intelligent work. As discussed in Chapter 5, the Agentic Genome Map exists to solve this problem. It provides a shared mental model for how autonomous capabilities should be structured so they can move from pilots to platforms. We apply this mapping to enterprise functions in this chapter.

In this chapter, we explore agentic workflows for:

- Executive and Board Agents
- Finance and Audit Agents
- HR and Talent Agents
- Sales and Marketing Agents
- IT, Security, and Compliance Agents
- Operations, Supply Chain, and Facilities Agents
- Legal, Risk, and Records Agents

A cross-cutting capability underpins every function described below: context-as-a-service. Agents across finance, HR, supply chain, IT, and legal all require governed access to enterprise content—documents from multiple repositories, B2B transaction data, and IT operations telemetry. This represents a fundamental shift from enterprise content management (storing and governing documents) to enterprise content delivery for agentic consumption (serving governed content to agents on demand, with access controls, compliance metadata, audit trails, and sovereignty routing embedded in every delivery). In every workflow that follows, the quality of agent reasoning depends directly on the quality of the governed context it receives.

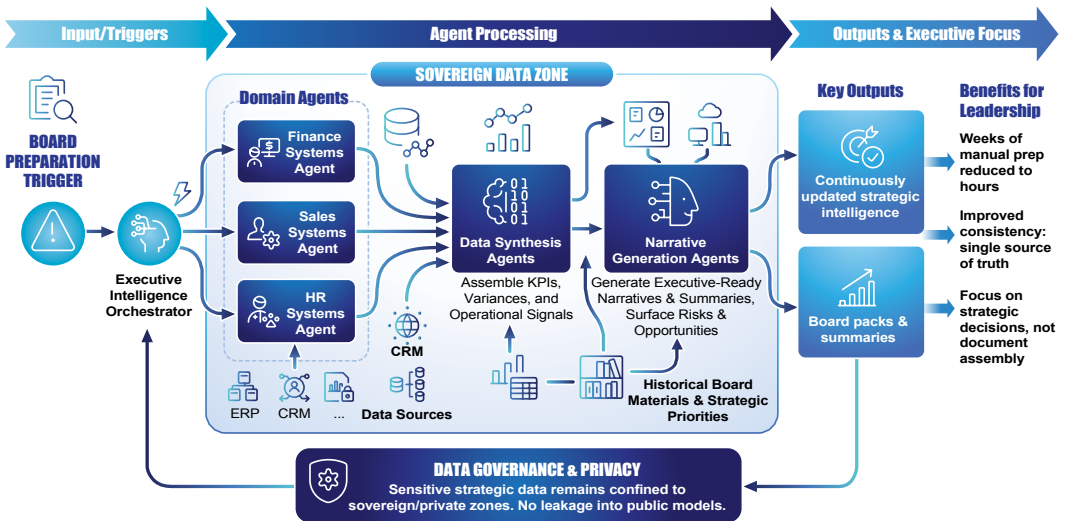
Executive and Board Agents

Goal: Shift from reactive reporting to predictive strategic navigation.

Agentic capabilities at the executive level transform leadership from consuming static reports to navigating continuously updated strategic intelligence. By automating the data gathering and synthesis phases, the C-suite and Board can focus their limited time on high-stakes decision-making rather than data validation.

In board preparation, an **executive intelligence orchestrator** can trigger **domain agents** across finance, sales, and HR systems ahead of meetings. **Data synthesis agents** assemble KPIs, variances, and operational signals from ERP and CRM platforms, while **narrative agents** trained on historical board materials and strategic priorities generate executive-ready summaries that surface emerging risks and opportunities. This compresses weeks of manual preparation into hours, improves consistency through a single source of truth, and frees leadership to focus on decisions rather than document assembly. Sensitive strategic data remains confined to sovereign or private zones to prevent leakage into public models.

Agentic AI Workflow: Board Preparation & Strategic Intelligence



Beyond reporting, agentic systems enable continuous competitive war-gaming and scenario planning, as broken down in the table below.

Continuous Strategic War-Gaming Workflow

Component	Intelligence Layer	Role
Strategy Orchestrator	Orchestration & Governance	Acts as the "Brain," initiating simulations and coordinating specialized agents to ensure alignment with corporate goals.
Market Intelligence Agents	Data Acquisition & Signal Detection	Continuously scans external environments for competitor moves, AI legislation, and macroeconomic shifts.

Component	Intelligence Layer	Role
Simulation Agents	Predictive Modeling & Analytics	Runs complex "what-if" models to predict impacts on financial performance, P&L, and global supply chain stability.
Scenario Planning Engine	Synthesis & Reasoning	Synthesizes agent outputs into coherent strategic options, allowing leadership to visualize various "future states."
Strategic Resilience Layer	Executive Decision Support	Translates raw simulation data into actionable insights, shifting the focus from "what happened" to "how do we pivot."

Key Advantages:

- **Dynamic Resilience:** Moves beyond static annual plans to a living strategy that adapts to regulatory or market shocks.
- **Regulatory Readiness:** Specifically monitors emerging AI legislation to ensure compliance-by-design during strategic shifts. The competitive advantage here is created by a governed intelligence layer which ensures strategic decisions are grounded in trusted data, aligned with policy, and fully auditable under leadership oversight.
- **Reduced Latency:** Eliminates the "insight gap" between a market event and a leadership response.

Governance agents further extend this model into continuous ESG (Environmental, Social, and Governance) and compliance oversight, scanning operational data against sustainability commitments and internal bylaws and escalating deviations in near real time. This reframes governance from retrospective assurance to live operational control, in accordance with the "four pillars of governance" (metadata, permissions, lifecycle, auditability).

Finance and Audit Agents

Goal: Move from retroactive accounting to continuous financial assurance.

In finance, agentic systems transform procure-to-pay processes from manual reconciliation into continuous, autonomous flows. By automating the standard reconciliation path and intelligently managing exceptions, the organization can capture discounts and optimize working capital without increasing headcount.

Autonomous Accounts Payable and Invoice Reconciliation Workflow

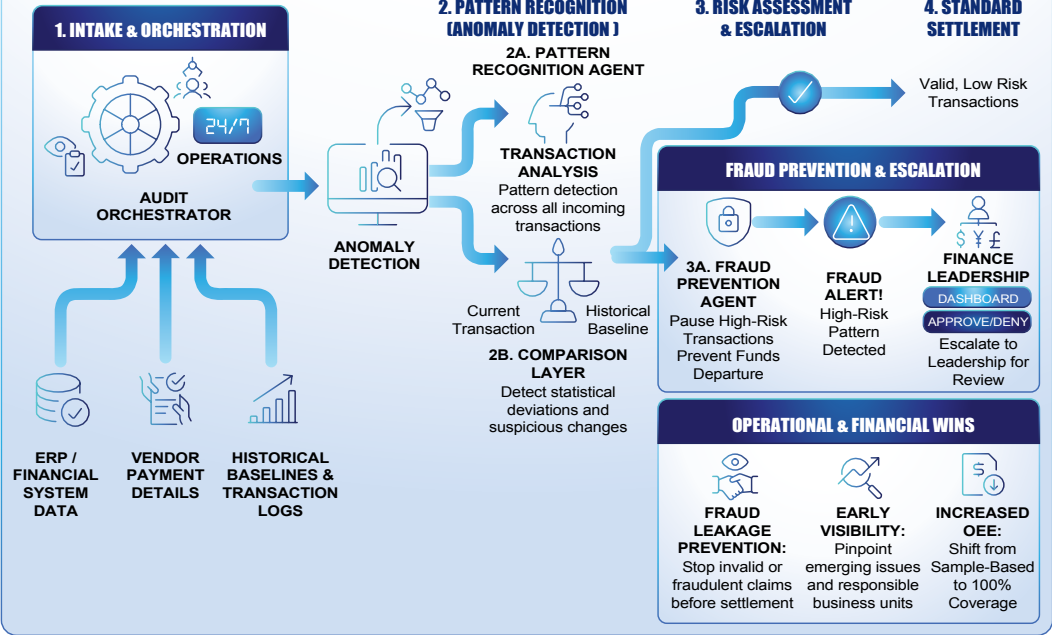
Component	Role	Intelligence Layer
Payables Orchestrator	End-to-End Flow Governance	Acts as the Central Controller, managing the lifecycle of an invoice from arrival to settlement and ensuring synchronization between procurement and finance systems.
Reconciliation Agent	Tri-Way Matching & Extraction	Employs Advanced OCR & Computer Vision to extract data, then applies logic to match the invoice against Purchase Orders (PO) and Goods-Received Records.
Resolution Agent	Exception & Conflict Management	A conversational and reasoning engine that autonomously engages suppliers or internal buyers to resolve price discrepancies or missing documentation.
Payment Agent	Cash Flow Optimization	Utilizes operational logic to schedule payments at the mathematically optimal time—balancing early-payment discounts against current liquidity needs.

Strategic Financial Wins:

- **Captured Margin:** By processing invoices in hours rather than weeks, the system never misses an early-payment discount (e.g., 2/10 Net 30), which can represent millions in annual savings for large enterprises.
- **Fraud & Duplicate Prevention:** The reconciliation agent performs instantaneous checks across the entire historical database, identifying duplicate invoices or suspicious “phantom” billing that manual reviews often overlook. The benefit is continued financial assurance, where every transaction is verified, every anomaly surfaced, and every decision traceable to governed data.
- **Shift to Strategic Finance:** With labor-intensive manual entry eliminated, the finance team can shift their focus to higher-value tasks like strategic sourcing, treasury management, and vendor negotiations.

With agentic AI, audit and fraud detection become continuous rather than periodic.

AGENTIC AI WORKFLOW: CONTINUOUS AUDIT & 100% TRANSACTION COVERAGE



In this workflow, an **audit orchestrator** can operate around the clock, with **pattern recognition agents** analyzing transactions against historical baselines to detect anomalies such as sudden changes in vendor payment details (such as a vendor changing bank details just before a large payment). When high-risk patterns emerge, **fraud prevention agents** can pause transactions and escalate to finance leadership before funds leave the organization, evolving audit from “sample-based” to “100% coverage.”

Treasury functions similarly benefit from dynamic cash-flow forecasting, where a **treasury orchestrator** aggregates real-time data from sales pipelines, payables, and macroeconomic data to model liquidity scenarios across near-term horizons. A **forecasting agent** runs probabilistic models to predict liquidity positions for the next 30/60/90 days. The result is a shift from backward-looking reporting to forward-looking financial steering.

As Michelin’s global transformation proves below, the agentic enterprise is not built in a vacuum. It is built on a foundation of governed, cloud-based integration that turns every invoice, shipment notice, and payment into a machine-readable signal—allowing the organization to finally trade manual oversight for continuous, autonomous financial assurance.



Michelin



We were looking for a scalable solution with security. We found that B2B Managed Services was completely in line with our strategy and this is why we moved to Managed Services. I would say that we are extremely happy.

— EDI Manager, Michelin

Michelin is one of the two largest tire manufacturers in the world. The company is dedicated to sustainably improving the mobility of goods and people by manufacturing and marketing tires for every type of vehicle, including airplanes, automobiles, bicycles/motorcycles, earthmovers, farm equipment, and trucks. Headquartered in Clermont-Ferrand, France, Michelin is present in more than 170 countries, has 111,200 employees and operates 67 production plants in 17 different countries.

The dynamic nature of the automotive industry and the significant growth opportunities in emerging markets presents a significant challenge to companies like Michelin. In order to minimize their supply chain risk, moving production capacity to other countries, or working more closely with customers in emerging markets, the company required a flexible B2B infrastructure.

A cloud-based integration platform supports Michelin's internal, financial operations: Trading Grid provides a secure, scalable, highly reliable platform for the exchange of transactions, such as purchase orders, shipment notices, commercial invoices, and payment instructions with its business partners. B2B Managed Services offers Michelin global customer support to proactively manage and monitor all B2B processes and business transactions, ensuring their B2B operations are running smoothly around the clock. Active Invoices enables Michelin to receive invoices electronically from their suppliers—and to send invoices electronically to their customers.

HR and Talent Agents

Goal: Deliver hyper-personalized employee experiences at scale.

In HR, agentic orchestration enables end-to-end employee journeys to be managed as coordinated workflows rather than disconnected tasks. This agentic workflow redefines the employee experience by shifting from a series of disconnected tickets to a unified, automated journey. By synchronizing IT, Facilities, and Learning departments, the organization ensures that every new hire is fully equipped and engaged from the moment they sign their contract, significantly reducing the “time-to-productivity.”

Intelligent Employee Onboarding and Journey Orchestration Workflow

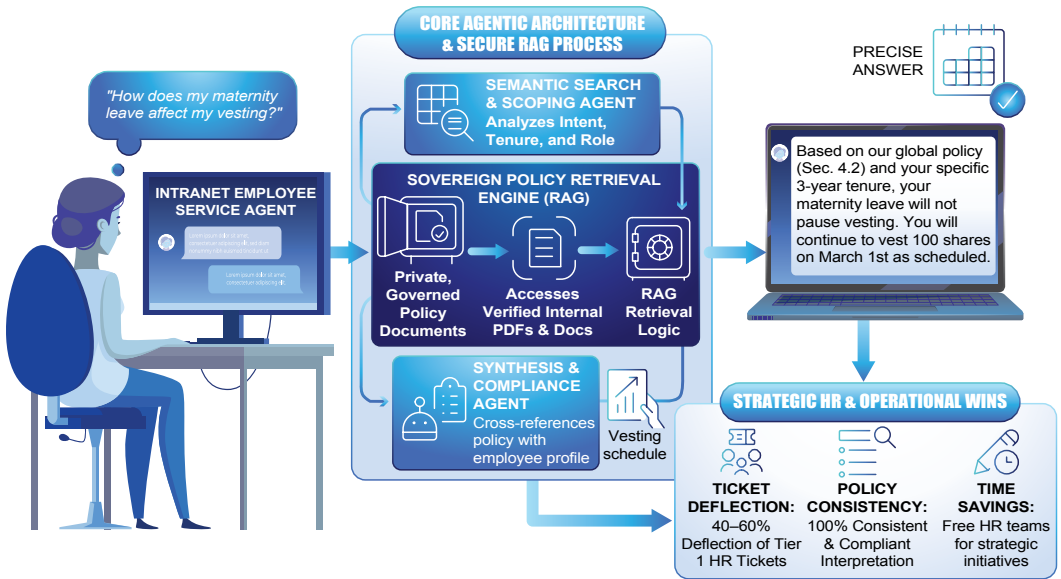
Component	Role	Intelligence Layer
Onboarding Orchestrator	Journey Lifecycle Governance	Acts as the Central Command, triggered by a signed contract to synchronize the sequencing of IT, Facilities, and HR tasks.
IT Provisioning Agent	Secure Access & Hardware Setup	Utilizes policy-based logic to autonomously order hardware, create system credentials, and assign software licenses based on the employee’s role.
Facilities Agent	Physical Workspace Allocation	A resource management engine that queries office floor plans and asset logs to assign a desk, locker, or parking pass.
Learning & Engagement Agent	Personalized Knowledge Transfer	Employs Natural Language Understanding (NLU) to guide new hires through compliance docs and schedule role-specific training modules via a conversational interface.

Strategic HR Wins:

- **Day One Productivity:** By automating the provisioning lag, employees arrive on day one with their laptop, badges, and system access already active, rather than spending their first week waiting for tickets to close.
- **Radical Administrative Relief:** HR Business Partners (HRBPs) are freed from the manual checklist drudgery of emailing different departments, allowing them to focus on high-value culture building and talent strategy.
- **Compliance With Policy and Regulations:** Across HR and talent functions, the advantage is not only agents coordinating employee journeys, but a governed workforce intelligence that ensures every policy decision, access change, and talent movement remains consistent, transparent, and accountable.

In our second agentic workflow scenario, an employee service agent sits within the intranet. Employees ask complex questions (“How does my maternity leave affect my vesting?”).

Agentic AI Workflow: Corporate Employee Intranet



The agent pulls from private, governed policy documents to give a precise answer, rather than a generic link. Using RAG to securely access human generated content, this can deflect 40-60% of Tier 1 HR tickets while ensuring consistent policy interpretation.

A **talent orchestrator** extends agentic capabilities into skills gap analysis and internal mobility, mapping workforce capabilities against strategic needs. A **career path agent** suggests internal roles or learning paths to employees based on their profile and company priorities. Recruitment costs are reduced by reframing talent management as a continuous optimization loop rather than a periodic workforce planning exercise.

In a sovereign HR environment, orchestration is the only way to ensure that strict regulatory deadlines—like the 60-day threshold for civil service appeals—are met without massive human oversight. By deploying eAppeals, the County of Los Angeles (below) has moved beyond fragmented legacy systems to a coordinated workflow that prioritizes tasks and manages cycle times, mirroring the onboarding orchestrator model to protect both the employee experience and the organization's legal standing.

County of Los Angeles, Department of Human Resources



The system's ease of use and the level of automation means the County will expend less human capital on the mundane process of the appeal and more human capital on the actual analysis and customer service of the appeal.

– Department CIO, County of Los Angeles Department of Human Resources

As the most populous county in the United States, the County of Los Angeles manages human resources for more than 250,000 applicants and over 110,000 employees across all departments. Its Appeals Program provides independent review of protests involving exam related issues, as well as protests involving discipline or personnel actions. The County receives close to 5,500 appeals every year. HR professionals in the Department of Human Resources (DHR) Appeals Unit review every submission, determining its merits in accordance with county and departmental policies, Civil Service Rules, and other standards.

Though DHR Appeals staff members worked as quickly as possible using former methods, any appeal exceeding a 60-day threshold may go to the Civil Service Commission (CSC), an administrative jurisdictional hearing body. The escalation demanded more people, oversight, and expense. The County's system lacked the ability to track progress or to produce an audit trail; appellants could only check the status of an appeal by contacting DHR professionals directly, which resulted in customer service complaints. The system also had limited security controls and did not integrate with other county applications. LA County wanted to offer online appeals submission with an automated interface to support county employees in the 60-day turnaround.

Working with trusted solutions provider, the County released eAppeals, an online appeals management system. The interface is automated: A person submits an appeal online and it is automatically delivered to the HR team, sending direct communication to the County department. Skipping the previous email and phone chase, departments then upload necessary documentation directly to the eAppeals system. It allows HR managers to track tasks, workload, and case cycle times to ensure action is prompt and in compliance with the 60-day deadline. Unlike the previous appeals system, eAppeals provides an audit trail: records show recipients, actions, and more with time stamps. Furthermore, the system provides effective controls to manage user access and passwords for confidential cases. The system's ease of use and the level of automation means the County will expend less human capital on the mundane process of the appeal and more human capital on analysis and customer service.

Sales and Marketing Agents

Goal: Create a "segment of one" at enterprise scale.

Agentic systems enable marketing to move from campaign management to hyper-personalized journey orchestration. By automating the loop between content creation and financial allocation, the system ensures that every dollar of spend is dynamically routed toward the highest-performing micro-segments.

Adaptive Marketing and Campaign Optimization Workflow

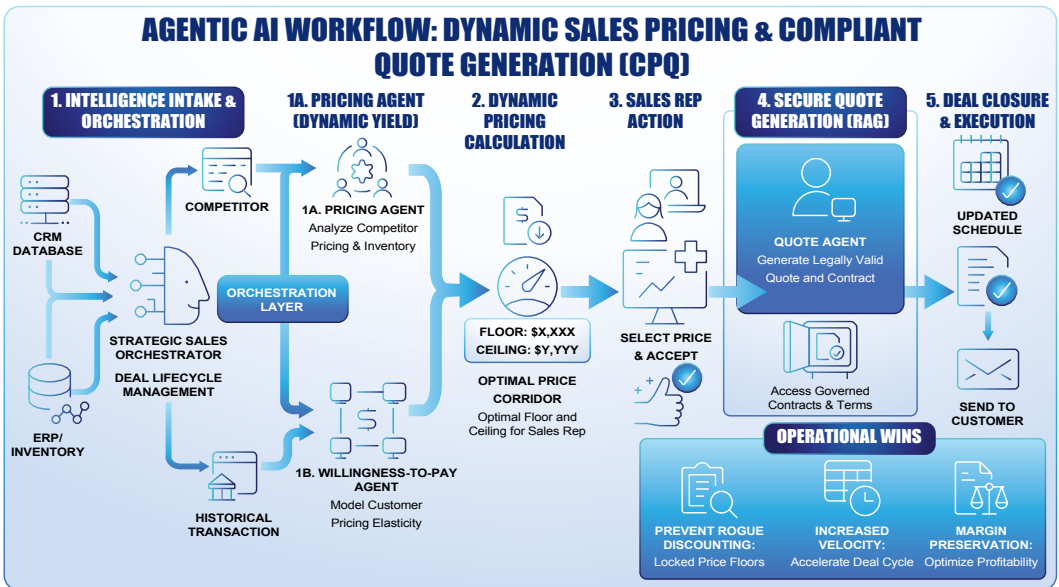
Component	Role	Intelligence Layer
Marketing Orchestrator	Campaign Strategy & Guardrails	Acts as the Central Strategy Hub, ensuring that all agents operate within the predefined brand voice, budget limits, and CRM segments.
Content Agent	Tailored Messaging Generation	A GenAI (LLM) engine that uses CRM data to create personalized email, social, and web copy specifically tuned for microscopic audience segments.
Deployment Agent	Omni-Channel Execution	An automation logic layer that manages the timing and delivery of content across fragmented platforms (Meta, Google, Email) to maximize open rates.
Analytics Agent	Performance Monitoring & Spend Pivot	A predictive feedback loop that analyzes real-time conversion data and automatically reallocates budget from underperforming ads to high-yield segments.

Strategic Marketing Wins:

- **Hyper-Personalization at Scale:** Traditional marketing teams often struggle to personalize content. This workflow allows for thousands of unique messages tailored to individual behaviors without increasing headcount.
- **Continuous CAC Optimization:** Instead of waiting for a weekly report to stop a failing ad, the analytics agent identifies spend waste in minutes, lowering the Customer Acquisition Cost (CAC) through instant reallocation.

- Brand Consistency and Governance:** The marketing orchestrator serves as a brand vault, ensuring that even when the Content Agent generates 5,000 variations of an email, the core messaging remains compliant and on brand. A governed engagement framework ensures that every message, price, and customer interaction remains compliant with brand standards, budget controls, and approved commercial policies.

Sales processes similarly benefit from dynamic pricing and quote generation (CPQ). In this scenario, a **pricing agent** analyzes inventory, competitor pricing, and customer willingness-to-pay in real time to guide pricing decisions—giving the sales rep an optimal price floor and ceiling. **Quote agents** can then generate compliant contracts, reducing cycle time and preventing unauthorized discounting. This maximizes margin, speeds up deal velocity, and prevents rogue discounting.



At the top of the funnel, autonomous **sales development agents** engage inbound leads around the clock, qualifying intent and scheduling qualified prospects with human sellers. This shifts human sales effort toward high-value closing activities while maintaining continuous lead engagement.

The adaptive marketing workflow relies on a continuous feedback loop where analytics agents trigger immediate shifts in resource allocation. This move from static reporting to autonomous refinement is precisely what Sky IT Group (below) enabled for global fashion brands; by deploying a high-performance cloud platform that processes disparate retail signals, they allow brands to pivot inventory and pricing strategies in real-time, effectively serving as the intelligence layer for a 28% increase in client revenue.

Sky IT Group



Sky IT Group is a leader in global sell-through data collection, validation, and analytics. Its SKYPAD web-based reporting suite integrates advanced data collection and integration methodologies with an intuitive user-focused interface, giving merchandising, planning, and sales teams access to a powerful self-serve reporting platform.

As the company's client roster grew to over 50 global brands, including Theory, Alice & Olivia, Lacoste, Fendi, and Marc Jacobs, several challenges emerged around data variety, velocity, and volume. Specifically: *How can these high-profile brands collect critical data, run high-performance analytics, and visualize results to understand consumer behavior?*

Sky IT Group chose an Analytics Platform with data visualization to accelerate and refine data analysis, running in a cloud-based environment. Fashion brands use the solution to view and analyze data collected from multiple retail channels and disparate data sources. Retail clients have improved their ability to understand customer trends and behaviors, enabling them to fine-tune inventory placement, driving sales and margin increases, and reducing the risk that products need to be marked down.

Supplier collaboration has also improved, and retailers can allocate resources more precisely to ensure they maximize return and minimize overhead. Inventory is allocated based on where specific items are selling, placing products in locations with the highest demand. Clients gain visibility into weekly sales data with minimal lag between actual sales and data availability. They can easily track what's selling, in which stores, and at what quantities and price points, among other metrics. Stores are less likely to be overstocked with unpopular items and out-of-stock on popular items. This not only maximizes sales, but enhances brand loyalty, and reduces risk that products will need to be marked down, which has improved profitability. Sky IT Group clients have reported an average revenue increase of 28%.

Key operational improvements in this case study center on precision, responsiveness, and visibility. Organizations who make such shifts lay the foundations for agentic AI. By leveraging rapid, cloud-based intelligence loops, organizations can allocate resources more strategically—strengthening supplier collaboration while optimizing returns. Inventory management becomes demand-driven rather than forecast-bound, with adaptive allocation placing the right products in the right locations based on real-time localized buying signals.

Our clients have reported an average revenue increase of 28 percent. Better analytics and actionable intelligence translates into better retail decisions and higher profitability for our clients.

– VP Business Development, Sky IT Group

IT, Security, and Compliance Agents

Goal: Self-healing infrastructure and proactive defense.

This agentic workflow shifts IT operations from reactive “firefighting” to autonomous self-healing infrastructure. By automating the diagnostic and remediation loops, the system eliminates the war room delays typical of manual incident response, protecting service-level agreements (SLAs) without human intervention.

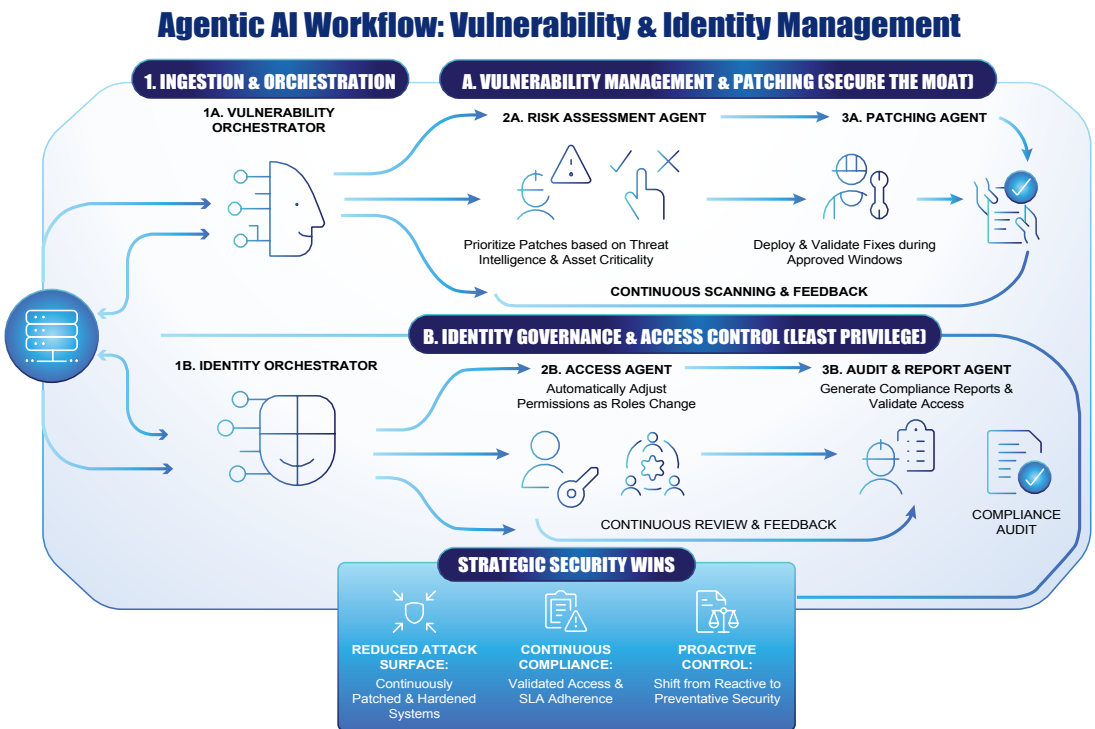
Autonomous IT Operations and Incident Remediation Workflow

Component	Role	Intelligence Layer
Incident Orchestrator	Global System Monitoring	Acts as the Central Watchtower, continuously ingesting telemetry data to detect performance anomalies that deviate from established “golden signals.”
Diagnostic Agent	Cross-Layer Root Cause Analysis	A pattern recognition engine that performs “log-stitching” across servers, networks, and apps to isolate root causes, such as the specific origin of a latency spike.
Remediation Agent	Automated Corrective Action	Employs deterministic logic to execute pre-approved safety maneuvers—such as container restarts, traffic rerouting, or code rollbacks—based on the diagnostic output.
Post-Mortem Agent	Knowledge Base Augmentation	A reasoning layer that summarizes the incident, updates internal documentation, and suggests permanent infrastructure hardening to prevent recurrence.

Strategic IT Wins:

- **Radical MTTR Reduction:** Bypassing the manual triage phase compresses the Mean Time to Resolution (MTTR) from hours of human investigation to seconds of machine execution.
- **Cascading Outage Prevention:** The speed of the Remediation Agent allows the system to quarantine a failing component before the error propagates through the rest of the microservices architecture.
- **Cognitive Load Shift:** Site Reliability Engineers (SREs) are freed from routine on-call alerts, allowing them to focus on high-value architecture design and proactive security posture.

Security functions benefit from **vulnerability and identity management**, as illustrated in the following diagram.



Vulnerability orchestrators coordinate risk assessment agents that prioritize patches based on threat intelligence and asset criticality, while **patching agents** deploy and validate fixes during approved windows. **Identity orchestrators** enforce least-privilege access, with **access agents** automatically adjusting permissions as roles change and **audit and report agents** generate audit reports for compliance reviews. Together, these capabilities reduce the attack surface, ensure compliance with SLAs, and shift security from reactive response to proactive, continuous control to “secure the moat.”

When agents act in production at regulated enterprises, every action creates compliance obligations. *What did the agent access? Was it authorized? What decision did it make and why? Can the reasoning chain be reconstructed?*

For agent compliance and cross-platform governance, agent **audit trail orchestrators** capture activity logs from every platform—Microsoft Copilot Studio, Salesforce Agentforce, Google Vertex AI Agent Builder, IBM watsonx Orchestrate, ServiceNow AI Agent Studio, etc.—and write them into records management with classification, retention, and legal hold applied automatically.

This cross-platform imperative is structural: single-platform governance is a feature; cross-platform governance—one policy, one console, one audit trail spanning all platforms simultaneously—is a product category. No platform vendor will ever index their competitors’ agent data. Meanwhile, agent registries provide automated discovery and cataloging of every agent deployed across the enterprise.

Shadow AI is proliferating like shadow IT did a decade ago; most CIOs cannot answer the question: *how many agents are currently deployed?* A compliance gap analysis might reveal: “247 agents deployed, 89 unregistered, 34 high-risk with incomplete governance”—the kind of visibility that drives immediate executive action. Agent FinOps capabilities add per-agent cost tracking, addressing Gartner’s prediction that 40% of agent projects will be canceled by 2027 due to 20–30x token cost overruns.¹² The EU AI Act high-risk rules, effective August 2026, make governance capabilities mandatory for any organization operating agents in regulated domains.

The real IT advantage is not only an infrastructure that heals itself faster, but also an operational environment where every automated action—every restart, patch, access change, and remediation—is executed within defined policies and preserved in a verifiable audit trail.

As the following case study demonstrates, the road to a self-healing infrastructure begins with total visibility. By turning manual security checks into a centralized, automated service, they have not only secured their cloud migration but have also built the watchtower necessary for autonomous agents to eventually diagnose and remediate threats in real-time.



A Global Beverage Bottling and Distribution Enterprise



A global beverage bottling and distribution enterprise operates dozens of production facilities and serves an extensive network of distribution centers across multiple countries. As the organization expanded its digital footprint, a transition to cloud-hosted applications introduced new complexity in managing application security at scale. This shift created an urgent need for a more structured approach to identifying vulnerabilities before they could impact operations.

As awareness grew around the potential consequences of a security breach in a cloud environment, the organization prioritized a solution capable of continuously evaluating application security across a diverse and distributed portfolio. The objective extended beyond detection to establishing a consistent, policy-driven framework that could standardize how applications—both internally developed and partner-provided—were assessed, governed, and brought into production.

Application Security as a Service (ASaaS) was implemented as a centralized capability, allowing the security team to scan applications across the portfolio in parallel. The service provides a streamlined method for assessing new applications against defined security standards prior to production deployment. A unified interface delivers end-to-end visibility into scanning activities, findings, and remediation priorities, effectively transforming security signals into actionable workflows that can be coordinated across teams and systems. By converting manual, point-in-time security checks into a continuously operating, orchestrated service, the organization established a scalable foundation for enterprise-wide application security. This model supports broader business objectives like regulatory compliance while introducing the architectural conditions for more advanced automation—where detection, prioritization, and response can increasingly be coordinated as part of an integrated execution layer, reducing latency between insight and action and setting the stage for progressively autonomous security operations.

Operations, Supply Chain, and Facilities Agents

Goal: Resilience and predictive optimization.

Agentic capabilities in operations enable predictive maintenance for critical assets. By closing the loop between real-time physics (Internet of Things [IoT] signals) and supply chain actions (parts and labor), the system ensures that critical equipment never reaches the point of catastrophic failure.

Predictive Maintenance and Asset Orchestration Workflow

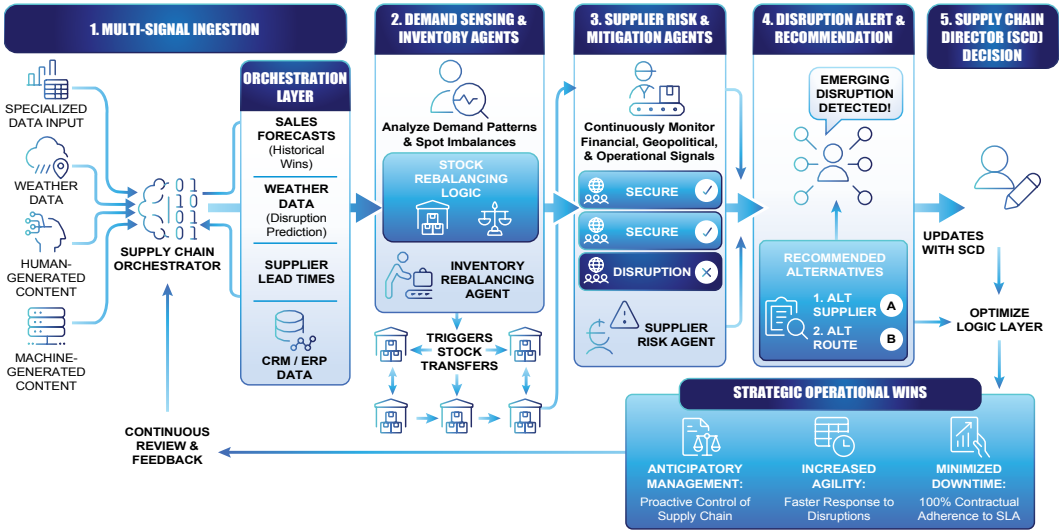
Component	Role	Intelligence Layer
Maintenance Orchestrator	Asset Lifecycle Governance	Acts as the Central Command, managing the integration between IoT telemetry, ERP inventory, and field service scheduling.
Prediction Agent	Anomaly & Failure Forecasting	A Deep Learning engine that analyzes vibration, acoustics, and heat signatures to identify “invisible” degradation patterns before they manifest as downtime.
Service Agent	Logistics & Resource Planning	A transactional logic layer that autonomously cross-references parts inventory, places orders, and matches technician skill sets with available service windows.
Optimization Agent	Lifecycle Strategy	A reasoning layer that evaluates the cost-benefit of a repair versus a full replacement, aiming to maximize the Total Cost of Ownership (TCO).

Strategic Operational Wins:

- **Guaranteed Uptime:** By resolving issues during planned windows, the organization eliminates the high costs of emergency downtime and lost production hours.
- **Just-In-Time Parts Management:** The service agent ensures that capital isn’t tied up in just-in-case spare parts inventory; instead, parts are ordered exactly when the prediction agent forecasts a need.
- **Extended Asset Longevity:** Continuous, minor interventions based on actual wear patterns prevent the cascading damage that typically occurs during a major mechanical failure.

Supply chain orchestration extends this model to demand-sensing and inventory balancing. **Supply chain orchestrators** integrate sales forecasts, weather data, and supplier lead times, while **inventory agents** rebalance stock across locations to prevent stockouts or overstocking. **Supplier risk agents** continuously monitor financial, geopolitical, and operational signals. Orchestrators alert the Supply Chain Director to emerging disruptions and recommend pre-approved alternatives. The result is a shift from reactive logistics to anticipatory supply chain management, as illustrated in the following diagram.

Agentic AI Workflow: Supply Chain Orchestration



Beyond sensing and recommending, agentic supply chain capabilities are evolving toward **autonomous transaction execution**. AI agents will potentially intermeditate \$15 trillion in B2B purchases by 2028—commerce that flows through trading networks.¹³ Agents are moving from reading transaction data (querying supplier performance, compliance exceptions) to writing transactions: placing purchase orders, onboarding suppliers, matching invoices, and enforcing trade compliance.

This progression must be designed for graduated autonomy with configurable human-in-the-loop thresholds mirroring existing procurement delegation authority—auto-approve routine transactions below defined thresholds require human approval. Critically, agent-executed transactions must carry the same governance, non-repudiation, and auditability as human-executed transactions. Trust is built by letting organizations see agent intelligence value (Phase 1) before asking them to trust agent execution authority (Phase 2).

Resilient operations emerge when supply-chain automation is paired with governance—ensuring that every maintenance intervention, inventory movement, and agent-executed transaction remains authorized, policy-aligned, and fully auditable.

The goal of agentic supply chain and operations is predictive optimization—ensuring assets perform at peak efficiency throughout their entire lifecycle. Philips Healthcare (below) demonstrates this maturity by leveraging a cloud-based analytics platform to move beyond reactive service models. By predicting potential system failures before they impact clinical operations, they have effectively synchronized physics-based signals with service actions to maximize return on investment and ensure high availability for healthcare providers worldwide.

Philips Healthcare



Our Analytics Database-powered predictive maintenance system, built on vast amounts of data and advanced AI models, allows us to detect and address potential issues before they impact clinical operations. This improves the reliability of our equipment and enhances patient outcomes and satisfaction.

– Mauro Barbieri, Principal Architect Service, Philips Healthcare

Philips Healthcare is a global leader in health technology, committed to improving lives through innovative solutions that enhance clinical performance and reduce the cost of ownership for medical imaging systems.

Philips Healthcare faced significant challenges with the maintenance of its advanced medical imaging systems, like MRI and CT scanners. These machines are essential for patient diagnosis and treatment, requiring high availability to ensure optimal clinical performance and predictable costs. However, unplanned downtime due to maintenance issues not only disrupted healthcare services but also posed a financial burden on healthcare providers. Recognizing the need for a more reliable and efficient service model, Philips Healthcare set out to shift from a reactive to a proactive maintenance approach.

Philips Healthcare leveraged a predictive maintenance system that uses AI-driven analytics to minimize equipment downtime, ensuring high availability and uninterrupted patient care. The AI-driven platform analyzes vast amounts of data collected from various medical devices to predict potential system failures before they occur. By utilizing AI, Philips Healthcare can process complex datasets efficiently and identify patterns that indicate imminent issues, allowing them to take preventive action well in advance. Millions of medical system log files are processed daily. This directly supports Philips Healthcare's predictive maintenance and has led to a 30% reduction in equipment downtime. This improvement ensures that critical medical imaging systems are more reliably available for patient care, reducing delays in diagnosis and treatment.

Legal, Risk, and Records Agents

Goal: Speed to contract and uncompromised compliance.

In legal functions, agentic orchestration accelerates contract lifecycle management. By moving the first pass of contract review from expensive external counsel to autonomous agents, the organization can uphold its risk thresholds at machine speed, ensuring that no deal is delayed by standard redlining.

Automated Contract Lifecycle and Legal Orchestration Workflow

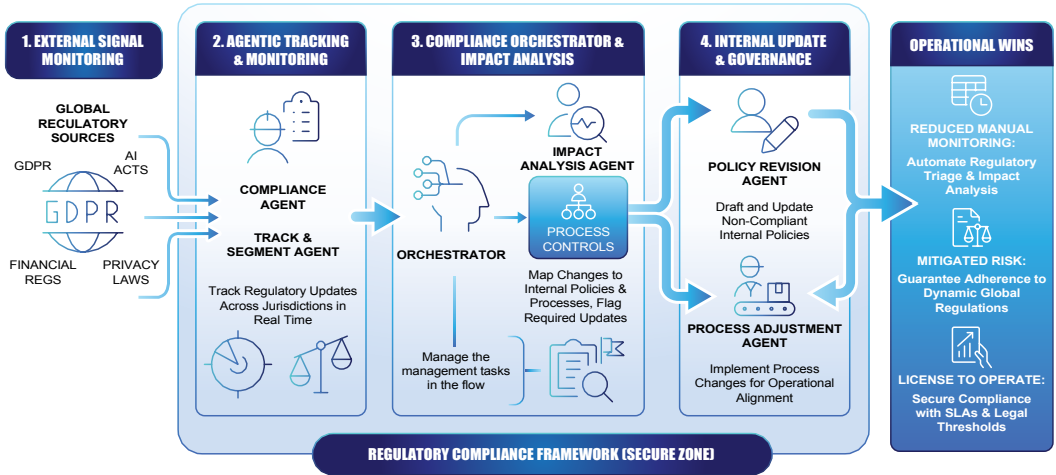
Component	Role	Intelligence Layer
Legal Orchestrator	Workflow & Lifecycle Governance	Acts as the Central Command, routing contracts between internal stakeholders, counterparties, and specialized review agents based on deal value or type.
Review Agent	Clause Analysis & Redlining	A semantic reasoning layer that compares third-party paper against internal "Gold Standard" playbooks to identify deviations in indemnity, liability, or jurisdiction.
Drafting Agent	Language Synthesis & Substitution	An LLM-driven generative layer that automatically replaces non-compliant clauses with pre-approved alternative language sourced from the corporate legal library.
Risk Compliance Agent	Threshold Monitoring	A deterministic logic gate that flags high-risk concessions for human General Counsel (GC) approval while green-lighting standard terms.

Strategic Legal Wins:

- **Accelerated Deal Velocity:** By automating the initial redlining phase, the legal review step in the sales cycle is compressed from weeks to minutes, preventing momentum loss during negotiations.
- **Significant Spend Reduction:** Organizations can dramatically reduce their reliance on external counsel for routine contract reviews, reserving expensive billable hours for high-stakes litigation or complex M&A.
- **100% Policy Adherence:** Human reviewers might miss a subtle change in an indemnity cap; an agentic system ensures that every contract, regardless of volume, is strictly cross-referenced against the latest board-approved risk thresholds. Across every legal function, the competitive advantage is not just the agent performing the task—it is the governed evidence trail proving the work was done correctly, within policy, and under appropriate oversight.

Regulatory change management also benefits from agentic monitoring. **Compliance agents** track regulatory updates across jurisdictions (such as changes to GDPR or new AI Acts), while **impact analysis agents** map these changes to internal policies and processes, flagging required updates. In this instance, agentic AI reduces manual compliance monitoring and risk, while protecting the organization’s license to operate.

Agentic AI Workflow: Regulatory Change Management & Compliance



In litigation contexts, **discovery agents** can scan millions of documents and emails across the enterprise, apply semantic search to identify relevant files, and assemble early case assessments for counsel, reducing discovery costs and time-to-insight.

The transition to an agentic enterprise is ultimately about converting massive data volume into immediate, actionable, and governed intelligence. For the national financial crimes investigation agency featured below, this would mean moving beyond the constraints of traditional, resource-intensive investigations to a scalable system capable of reviewing hundreds of thousands of documents daily. By automating the discovery and triage phases, the agency could drastically reduce its time-to-insight, ensuring that justice is neither delayed nor defeated by the sheer velocity of digital evidence.



Financial Crimes Investigation Agency



Modern economic crime investigations routinely involve hundreds of thousands of digital documents—from emails and financial records to communication logs—making traditional review processes slow and resource-intensive.

To address this challenge, let's assume a national financial crimes investigation agency deployed an AI-driven investigative platform that uses intelligent agents to ingest, classify, and analyze digital evidence. These agents automatically recognize patterns, group related documents, map relationships between individuals and organizations, and construct timelines of activity. By orchestrating multiple agents across the investigative workflow, the system prioritizes the most relevant evidence and enables analysts to focus on high-value investigative insights rather than manual document review.

The result is a scalable investigative environment capable of processing and prioritizing more than 100,000 documents per day. Instead of replacing investigators, the agentic system acts as a force multiplier—continuously scanning and organizing digital evidence while human experts interpret findings and build prosecutable cases. In an era where economic crime generates ever-growing volumes of digital data, this approach enables the agency to accelerate investigations and ensure that justice keeps pace with the scale of modern fraud.

EIM as the Agentic AI Backbone

As previously discussed, agentic AI involves goal-directed systems that can plan and execute tasks with a degree of autonomy rather than merely responding to prompts. In enterprise deployments, that autonomy must be anchored to governed data and controlled processes. This is where EIM platforms matter. They provide the institutional memory, evidentiary integrity, and policy context that allow agents to act in ways that are explainable, auditable, and compliant.

Without this backbone, agentic AI devolves into clever but unsafe automation. With it, autonomy becomes a managed capability that enterprises can trust, scale, and govern.

Here's a breakdown of agentic aspects of EIM functionality:

- **A Content Server** functions as the authoritative content substrate for standards, manuals, investigation reports, decision rationale, and briefings—the institutional memory layer that agents retrieve from and cite back into as evidence.
- **Archive Center** extends that memory into long-term retention and evidentiary integrity. In regulated industries and the public sector, retention, immutability, and defensible disposal are not optional records features; they form the bedrock of accountability.
- **Intelligent Capture** is an ingestion engine for the agentic enterprise: operator submissions, scanned records, evidence documents, and forms become structured, classified, and usable by downstream agents. Capture is where the physical world becomes machine-actionable policy evidence.
- **Knowledge Discovery** includes AI-powered products for data discovery, security, and intelligence. These tools automate data classification, risk management, and insights discovery across complex, hybrid environments—enriching metadata and transforming unstructured data into enterprise intelligence.
- **Analytics and Generative AI** provide the analytics and signal layer: safety trend signals, incident metrics, forecasting, and pattern detection that can trigger orchestrations and prioritize human attention. In an agentic deployment, analytics becomes a sensing mechanism feeding orchestrators, not just dashboards feeding humans.
- **Process Automation** anchors case-centric orchestration: certification cases, inspection cases, enforcement cases, and incident casefiles are exactly the kind of structured, repeatable, policy-driven workflows where agentic systems deliver value without becoming reckless. A low-code development platform used to build, deploy, and manage intelligent process automation and dynamic case management applications is a natural home for human-in-command gates because casework is already designed around approvals, escalations, and traceability.

- **Information Governance** provides the guardrails: retention schedules, legal holds, privacy controls, permissions, records declarations, and lifecycle management. In agentic systems, governance doesn't only oversee content—it regulates what agents are allowed to do with content, and what must be logged and retained.
- **Enterprise Service Management (ESM)** is a unified platform that automates services across IT, HR, facilities, and other departments using low-code configuration and embedded generative AI. Its AI capabilities streamline ticket intake, categorization, workflow routing, and self-service, reducing resolution times and operational friction. Over time, AI agents can learn from historical tickets, resolutions, and outcomes to autonomously resolve recurring issues and recommend next best actions. Together, these capabilities enable organizations to scale service delivery, improve employee experience, and move toward agentic, self-healing service workflows across the enterprise.
- **Business Networks** extend the enterprise's information substrate across external ecosystems by securely connecting partners, systems, and data to automate B2B integration, enable real-time visibility into transactions and supply chain flows, and underpin collaborative digital processes that feed trusted signals and partner context into agentic workflows.

Agentic AI is only as trustworthy as the enterprise memory it is grounded in. The important takeaway is that EIM platforms provide the institutional memory that makes autonomous action not only defensible but trusted to do the job it's given based on the appropriate unique business context.

Benefits of Agentic AI in the Enterprise

The pattern across executive, finance, HR, sales, IT, operations, and legal is consistent: agentic AI delivers value when it is designed as an operating model, not a collection of tools. Agents create leverage when they are grounded in enterprise memory, coordinated through orchestration, and governed with clear decision rights, auditability, and escalation paths. This is why the “agentic organization” is fundamentally about workflow redesign—replacing fallible handoffs with orchestrated, measurable outcomes—while keeping humans in command of policy and accountability.

Agentic AI improves resilience and adaptability by turning operations into learning loops. Responsible autonomy depends on structured, traceable data and clear policy constructs; when those foundations are in place, enterprises can link knowledge to situational context and continuously optimize performance.

At the same time, the path to value is not automatic. Hype-driven deployments, weak governance, and poor data foundations can derail pilot projects. The organizations that win will be the ones that treat agentic capabilities as enterprise infrastructure. They will standardize patterns, centralize oversight across the agent estate, and iterate continuously based on performance and risk signals.

Executive Implications

CAIO

Enterprise-scale autonomy requires consistent governance. The CAIO establishes the standards—risk classification, auditability, and human oversight—that ensure agents operate safely across functions and platforms.

CIO

Agentic scale depends on a shared platform rather than isolated solutions. The CIO stewards the enterprise agentic platform, providing orchestration, secure integration, and governed context that allow business units to deploy intelligent workflows safely.

CFO

Measurable value comes from shifting high-volume workflows (P2P, audit, pricing, maintenance) to continuous assurance and optimization, with exceptions routed to humans-in-command and ROI tracked by domain.

CHRO

Hyper-personalized employee experiences at scale require HR to operationalize human-agent collaboration, defining escalation paths, role boundaries, and stewardship for a blended workforce.

CDO

Enterprise-grade autonomy depends on EIM as institutional memory—metadata, permissions, lifecycle, and auditability—so agents can act on governed context rather than brittle, untraceable data flows.

COO

Cycle time, resilience, and service reliability improve when orchestration redesigns cross-functional work into measurable execution loops, turning operations from reactive handoffs into continuously optimized processes.

Now that we've examined agentic AI applied across enterprise functions, let's take a look at how agentic AI could be applied—according to the Agentic Genome Map—across organizations in both the private and public sectors.



Chapter Seven

Applying Agentic Capabilities to the Private and Public Sectors

This chapter illustrates what agentic AI transformation looks like in practice across industries in the private sector, as well as in public sector at the federal, state/provincial, and local/municipal levels. For each sector, we provide examples of specific workflows that agentic AI might orchestrate. These are end-to-end agentic workflows that link planning to execution, sensing to response, and policy to action across complex operational environments.

In the previous chapter, we showed how enterprises can streamline the flow of decisions and actions across systems, teams, and value chains by combining agents that reason, plan, and act with orchestration layers to coordinate execution under policy and human oversight. By applying the Agentic Genome to industries, organizations move from fragmented automation to orchestrated intelligence—embedding reasoning and execution across the value chain to include suppliers and partners in the supply chain.

In this chapter we examine the following Agentic AI Workflows by industry:

- **Banking and Finance (pg. 123):**
 - Financial Crime and Anti-Money Laundering
 - Loan Origination and Credit Underwriting
 - Autonomous Treasury Management and Liquidity Forecasting
- **Food and Beverage Manufacturing (pg. 127):**
 - Agentic Food Safety and Traceability Compliance
 - Intelligent Yield Optimization and Quality Control
 - Predictive Procurement and Production Planning
- **Transportation and Logistics (pg. 132):**
 - Dynamic Route Optimization and Disruption Management
 - Autonomous Predictive Fleet Maintenance
 - Agentic Freight Documentation and Customs Clearance
- **Automotive (pg. 136):**
 - Predictive Manufacturing Asset Maintenance and Quality Assurance
 - Intelligent Aftermarket Parts and Logistics Optimization
 - Agentic Warranty Claims Processing and Fraud Detection
- **Healthcare (pg. 140):**
 - Value-Based Care Performance Analytics
 - Predictive Medical Imaging Asset Maintenance
 - Agentic Patient Triage and Benefits Navigation
- **Pharmaceutical (pg. 145):**
 - Autonomous Regulatory Affairs and Submission Management
 - Agentic Drug Discovery and Clinical Trial Optimization
 - Intelligent Batch Record Review and Release

- **Process Manufacturing (pg. 150):**
 - Autonomous Batch Record Management and Compliance
 - Agentic Process Control and Energy Optimization
 - Predictive Asset Maintenance and Reliability
- **Oil and Gas (pg 155):**
 - Intelligent Production Optimization and Reservoir Management
 - Autonomous Asset Integrity and Predictive Maintenance
 - Automated Regulatory Compliance and HSE Monitoring
- **Energy and Utilities (pg. 159):**
 - Autonomous Grid Optimization and Load Balancing
 - Predictive Asset Maintenance for Utility Infrastructure
 - Automated Regulatory Compliance and Environmental Reporting
- **Public Sector – Federal (pg. 164):**
 - Autonomous Border Risk Assessment and Clearance
 - Agentic Benefits Eligibility and Case Management
 - Intelligent Federal Procurement and Contract Management
- **Public Sector – State/Provincial (pg. 169):**
 - Agentic Court Case and Docket Management
 - Autonomous Social Housing and Benefit Eligibility
 - Intelligent Service Request Classification and Resolution
- **Public Sector – Local/Municipal (pg. 173):**
 - Intelligent Citizen Service Resolution 311
 - Agentic Permitting and Zoning Compliance
 - Predictive Municipal Infrastructure Maintenance

The agentic shift is structural rather than incremental. End-to-end workflows become faster as decisions move closer to real-time execution, more accurate as agents reason over governed enterprise information instead of brittle point integrations, and safer as orchestration layers enforce policy, escalation paths, and human-in-command controls.

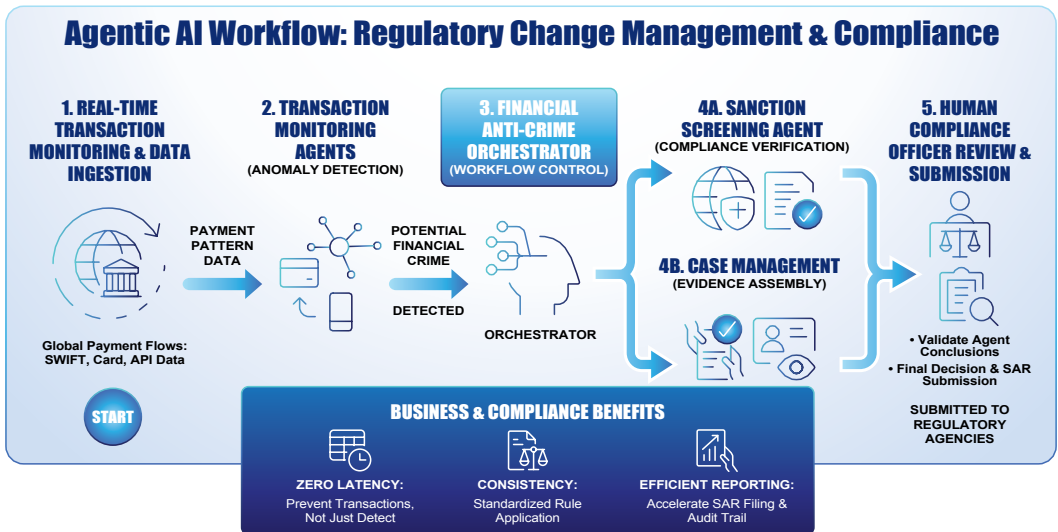
What follows is an exploration of how agentic AI, grounded in the Agentic Genome, transforms disconnected digital systems into coherent operating models that can sense, decide, and act across industry workflows at scale, with accountability and control.

Banking and Finance

Banking and financial institutions are operating in an environment defined by geopolitical volatility, fragmented regulation across jurisdictions, rising fraud and cyber risk, tighter capital requirements, and relentless pressure to improve customer experience while controlling costs. For European financial institutions, the regulatory timeline is already fixed. Beginning in August 2026, high-risk AI systems used in areas such as credit assessment, fraud monitoring, or financial decision-making must satisfy the EU AI Act's requirements for risk management, human oversight, auditability, and formal conformity assessment—or face penalties that can reach 7% of global revenue for non-compliance. Complementary frameworks including NIST AI RMF and ISO 42001 apply across non-EU jurisdictions. In this climate, agentic AI workflows shift financial operations from reactive control to continuous, intelligence-driven execution within a defensible governance framework.

Financial Crime and Anti-Money Laundering

Agentic AI equips financial institutions to conduct real-time risk management. Processes traditionally dependent on batch monitoring and manual review can be re-architected as orchestrated, always-on agentic AI workflows. In this workflow, a **financial anti-crime orchestrator** monitors transactions in real time across a global banking network, coordinating specialized **transaction monitoring agents** that analyze payment flows and use pattern recognition to detect anomalies indicative of money laundering or fraud. A **sanction screening agent** cross-references entities against global sanctions and watchlists (such as Office of Foreign Assets Control [USA] or the UN) to ensure compliance. When potential money laundering or fraud is detected, **case management agents** assemble the required evidence and generate a Suspicious Activity Report (SAR) for **compliance officers** to review and submit.



This approach can reduce false positives by up to 60%, lower the operational cost of compliance teams, and strengthen regulatory defensibility, all while protecting both customers and the institution from financial and reputational harm.

Loan Origination and Credit Underwriting

Agentic orchestration transforms document-heavy processes into responsive, end-to-end digital workflows.

Loan Origination and Credit Underwriting Workflow

Component	Role	Intelligence Layer
Lending Orchestrator	Process Control	The “Brain” that knows when to auto-approve, when to ask the applicant for more info, and when to call in a human expert.
Intake Agent	Data Capture	Uses OCR and NLP to pull data from IDs, paystubs, and forms; checks against “Know Your Customer” (KYC) policies.
Underwriting Agent	Analysis	Aggregates credit scores, real-time bank transaction history, and even alternative data sources.
Decision Agent	Risk Modeling	Applies risk models to determine loan eligibility and pricing. Runs the numbers against the bank’s risk appetite to set the interest rate and loan limit. Complex or high-risk cases are escalated to human underwriters for final judgment.

This workflow is a gamechanger for financial services because it moves paperwork at the speed of an API call. By the time a human underwriter even opens the file for a high-risk case, the agents have already done 90% of the legwork—gathering the data, spotting the red flags, and verifying the identity. Banks can reduce approval times from days to minutes while preserving rigorous credit standards and regulatory compliance, improving both customer experience and conversion rates.

Autonomous Treasury Management and Liquidity Forecasting

In this scenario, agentic capabilities shift financial steering from periodic reporting to continuous optimization.

Component	Role	Intelligence Layer
Treasury Orchestrator	Strategic Synthesis	Aggregates real-time data from enterprise financial systems and external market feeds. The “Conductor” that balances the need for yield against the necessity of liquidity, ensuring the enterprise never runs dry while maximizing ROI.
Cash Flow Agents	Predictive Modeling	Analyzes thousands of historical transaction patterns, payables, and receivables to model daily liquidity needs.
Risk Agents	Scenario Simulation	Runs simulations on market volatility (e.g., a shift in interest rates or sudden foreign exchange swings) to stress-test reserves.
Investment Agents	Yield Optimization	Evaluates short-term vehicle yields (Money Markets, Repo, etc.) to recommend allocation strategies to optimize returns on idle capital without compromising operational liquidity.

Instead of waiting for a monthly close to see cash positions, the CFO has a real-time dashboard of current and forecasted liquidity. By accurately predicting the floor of required cash, investment agents can put more capital to work in interest-bearing environments. The result is more accurate, real-time visibility into global cash positions, improved working capital efficiency, and reduced reliance on costly short-term borrowing—transforming treasury operations into a proactive, intelligence-driven function rather than a backward-looking control layer.

As financial institutions move toward agentic transaction execution, the design principle of graduated autonomy becomes essential. A rogue agent placing a \$10 million purchase order creates real financial exposure, and legal frameworks for agent-executed commercial transactions are still evolving. The prudent approach mirrors existing financial delegation authority: agents are granted configurable human-in-the-loop thresholds where routine transactions below defined limits proceed autonomously, while transactions above those limits require human approval. The institution first sees value from agent intelligence—analyzing transactions, identifying anomalies, recommending actions—before entrusting agents with execution authority. Trust is earned through demonstrated reliability, not assumed from capability.

True financial agility requires more than just faster math; it requires an orchestrator capable of synchronizing disparate data streams into a single source of truth. The success of Indian Ocean Bank (below) reinforces a fundamental truth of the agentic era: when a financial institution offloads the ‘drudgery of the desk’ to intelligent workflows, it doesn’t just save paper—it gains the agility to bridge continents and opportunities at the speed of the modern market.



Indian Ocean Bank



Our paper print has reduced by over 60 percent; from 6,000 pages per day to just 2,000. We also see clear improvements in resolution times. The teams can fulfil requests 35 percent faster than before because they don't need to change screens.

– Data and Integration Engineering Manager, Indian Ocean Bank

Indian Ocean Bank started with a single vision to connect people, places, and opportunities. Specializing in bridging Africa, Asia, and the world, it offers innovative financing solutions.

As a relatively young bank, serving more than 50,000 customers with its 400 employees, it has an inbuilt start-up mentality and the agility to embrace change when it is needed. The customer service team is dedicated to managing up to 1,500 daily customer queries, from balance queries to card passcodes, internet banking passwords, and payment requests. Towards the end of each month this number can often double, exponentially increasing the workload for the team. With a high level of manual effort and 6,000 pieces of paper moving between desks every day, this department was at risk of missing customer queries and getting overwhelmed during busy times.

To streamline customer service, Indian Ocean Bank deployed an AI- and automation-based solution, achieving 35% faster request fulfillment and a 60% paper reduction. The solution integrates with key systems in use around the bank and will route the ticket to the appropriate owner. When customers request a transfer from one account to another, it manages the checks involved in moving funds. The process will flag exceptions; for instance, if signatures don't match or when a client callback fails. The workflow will determine next steps in exception cases before submitting the payment and recording it in the document management system. AI capabilities go beyond automating routine tasks and can play a pivotal role in enhancing various aspects of the bank's IT Service Management (ITSM), customer service, security, and decision-making processes. This underscores the transformative impact of AI on the bank's operations, making it more agile, efficient, and customer centric.

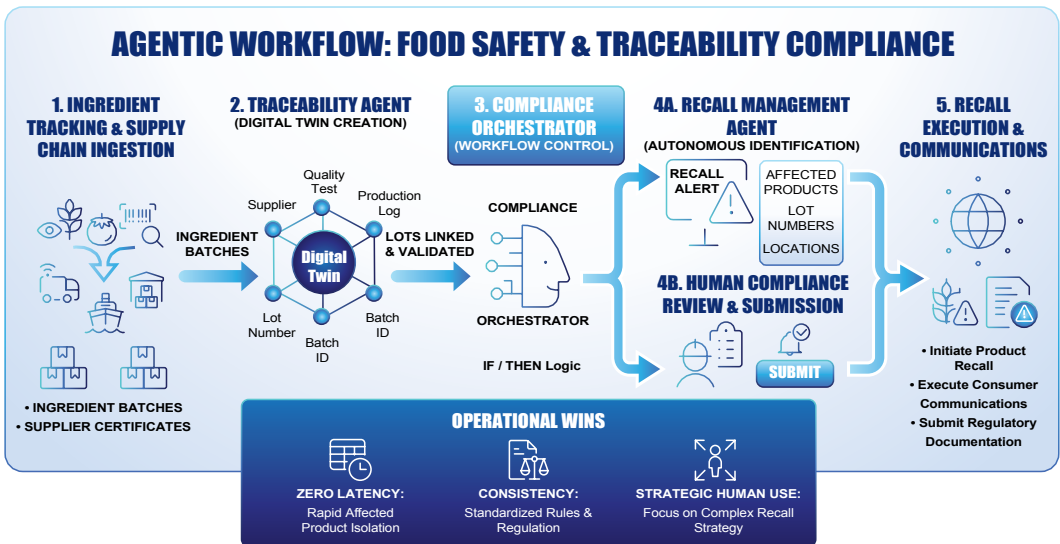
Food and Beverage Manufacturing

Food and beverage manufacturers face mounting pressure from fragmented data, siloed processes, and increasingly complex compliance and governance requirements across global supply chains. While IoT-enabled production lines generate vast operational data, it is often disconnected from quality, planning, and supplier systems, limiting real-time visibility and traceability across business trading networks. Bridging these gaps is driving a shift toward AI-powered automation that can connect data, orchestrate processes, and enable compliant, real-time optimization across the end-to-end manufacturing ecosystem.

Agentic Food Safety and Traceability Compliance

Agentic systems can be used to ensure food safety and regulatory compliance by automating traceability and recall readiness across the supply chain.

In this agentic workflow a **compliance orchestrator** tracks every ingredient batch from receipt through to final product, while a **traceability agent** continuously updates a digital twin of the supply chain, linking supplier certificates, quality test results, and production logs to specific lot numbers. If a safety alert or recall is triggered, a **recall management agent** can autonomously identify all affected products and their current locations—whether in warehouses, in transit, or at retail locations—and draft regulatory notifications and customer communications for human review.



This coordinated, agentic AI workflow reduces recall response times from days to minutes, minimizing public health risks and protecting brand reputation while ensuring compliance with regulations such as the Safe Food for Canadians Regulations (SFCR) and the FDA's Food Safety Modernization Act (FSMA). The approach directly supports governance and ethics requirements by embedding accountability and safety into operational workflows, and it leverages sovereign-zone architectural patterns to protect sensitive supply chain data and ensure compliance with regional regulatory requirements.

Intelligent Yield Optimization and Quality Control

In modern manufacturing environments, agentic systems can be applied to optimize yield and reduce waste by autonomously adjusting production parameters in real time. This shift moves manufacturing from Statistical Process Control (where you find out something is wrong after the batch is ruined) to Dynamic Autonomous Tuning (where the machine fixes itself on the fly).

Intelligent Yield and Quality Workflow

Component	Role	Intelligence Layer
Production Orchestrator	Global Optimization	Continuously ingests data from IoT sensors monitoring conditions such as temperature, humidity, and ingredient viscosity along the production line. Balances the Speed vs. Quality trade-off, ensuring adjustments in one stage don't create bottlenecks or waste downstream.
Quality Analysis Agent	Visual Inspection	Uses high speed computer vision and deep learning to identify microscopic defects or texture inconsistencies in real-time.
Process Adjustment Agent	Real-time Correction	Executes precise mechanical changes (e.g., increasing mixer speed by 2% or lowering oven temp by 5°C) to counteract detected drift.
IoT Sensor Grid	Environment Sensing	Captures high-frequency telemetry on physical variables like viscosity, moisture, and thermal gradients.

The "Zero-Waste" Advantage:

- **Non-Stop Production:** In traditional setups, a quality drift requires a "stop-and-fix" intervention. Here, the agent nudges the parameters back to center while the belt keeps moving.

- **Micro-Batch Precision:** This architecture allows for consistency even when raw material inputs vary (like different moisture levels in flour or different ambient humidity in the plant).
- **Reduced Scrap Rates:** By catching a deviation early, the agents prevent it from becoming a failure rate at the end of the line. Overall yield can be increased by 2-5%, for example, and scrap and waste by up to 20%.

This coordinated, agentic AI workflow demonstrates how agentic AI can be grounded in the operational technology layer of modern cloud architectures, connecting intelligent systems directly to physical machinery to enable closed-loop optimization on the factory floor.

Predictive Procurement and Production Planning

To maximize profitability, agentic systems can dynamically adjust procurement and production decisions based on shifting demand signals and volatile commodity prices. By connecting external market volatility directly to the factory floor, the organization can capture margin that is usually lost to lag or manual decision-making.

Predictive Procurement and Production Planning Workflow

Component	Role	Intelligence Layer
Supply Chain Orchestrator	Situational Awareness	Integrates sales forecasts, real-time commodity market prices, and weather patterns to maintain continuous situational awareness. Acts as the "Digital Command Center," ensuring that the procurement of raw materials and the production schedule are perfectly synchronized with the sales forecast.
Demand Sensing Agent	Market Analysis	Correlates non-traditional signals (e.g., changes in consumer preferences, local weather shifts, seasonal spikes) to predict SKU-level demand before it hits the order book.
Procurement Agent	Strategic Sourcing	Monitors commodity exchanges and uses Algorithmic Negotiation to execute buy orders at the "bottom of the curve" or secure future contracts when prices are trending up.
Production Scheduling Agent	Dynamic Rebalancing	Optimizes the manufacturing schedule in real-time, swapping low-margin production runs for high-demand, high-margin products based on the sensed forecast.



Key Competitive Advantages:

- **Margin Capture:** Purchasing raw materials at optimal price points and producing what the market *actually* wants in the moment significantly increases the “spread” per unit.
- **Agility at Scale:** Traditional supply chains take weeks to pivot. This agentic model can pivot in hours, allowing you to capitalize on a cold snap or a sudden competitor stock-out.
- **Inventory Optimization:** This coordinated approach can reduce raw material costs by 5–10% and inventory holding costs by 15–20%, while maintaining high service levels for key retail partners.

The workflow reflects agentic AI operating in a sovereign context, ensuring that sensitive financial data and partner information are handled securely while enabling real-time, automated decision-making across procurement and production planning.

To move from reactive logistics to predictive planning, an organization must first ensure its data sovereignty through a secure, cloud-based orchestration layer. The following feature reveals how one Farmer-Owned Cooperative replaced fragmented manual tasks with a governed platform, transforming document-intensive bottlenecks into high-velocity digital flows that are ready for real-time market adaptation.



A Farmer-Owned Cooperative



The global dairy cooperative processes millions of transactions annually. It was constrained by paper-heavy workflows and brittle on-premises systems. Critical operational documents—such as complex production and logistics manifests containing more than 150 data fields—were manually captured and re-entered into enterprise systems, slowing operations, increasing error risk, and tying up scarce IT resources. System reliability issues further limited the organization's ability to modernize core processes.

The organization modernized its content and process foundation by digitizing document flows and moving mission-critical systems to a cloud platform tightly integrated with ERP. This created the conditions for agentic AI workflows: process orchestrators that could coordinate end-to-end document lifecycles, with specialized agents to handle document ingestion, classification, data extraction, validation, and posting into core systems. By standardizing content pipelines and centralizing governance, the cooperative could establish a control plane where intelligent agents can operate safely within defined workflows and escalation paths.

Key Operational Improvements:

- **End-to-end document workflows are orchestrated** rather than manually coordinated, reducing re-keying, delays, and errors.
- **Intelligent capture agents could then extract structured data** from complex, multi-field documents and feed downstream systems automatically.
- **Centralized content services provide a governed "enterprise memory"** that agents can retrieve from and write back to with full auditability.
- **Orchestration enables new workflows**—such as maintenance and contract management—to be deployed rapidly without re-architecting core systems.

By shifting from fragmented automation to orchestrated workflows, the organization transformed document-intensive operations into scalable digital flows. The result is higher system reliability, faster execution, and a platform ready for agentic AI—where orchestrators manage end-to-end outcomes and agents perform specialized work under human-in-command governance.

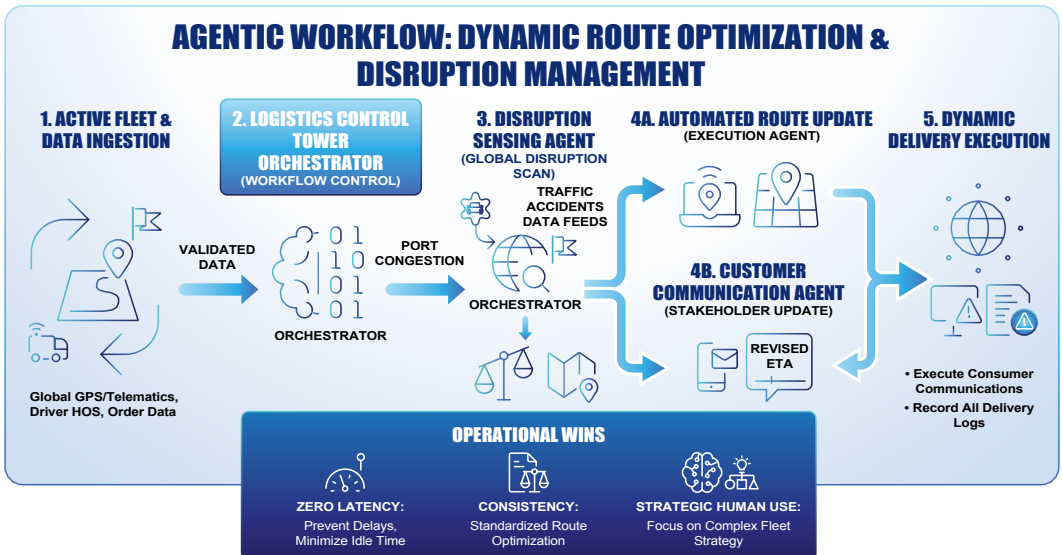
Transportation and Logistics

Transport systems sit at the intersection of critical infrastructure, regulation, economic activity, and public trust, coordinating the movement of goods and people across ports, borders, rail networks, highways, air corridors, and logistics hubs while operating continuously under growing regulatory and operational pressure.

As global supply chains become more interconnected, and disruption—everything from climate events to geopolitical shocks—propagates at machine speed, traditional digital transformation approaches are hitting their limits: data remains siloed across agencies and partners, coordination relies on fragmented systems and manual handoffs, and decision-making is largely reactive. The agentic enterprise model offers a different path forward by applying the Agentic Genome to transport and logistics, moving from disjointed automation to orchestrated intelligence that embeds reasoning and execution directly into the operating model of the transport ecosystem.

Dynamic Route Optimization and Disruption Management

In an agentic logistics enterprise, route planning evolves from static optimization to continuous orchestration, minimizing fuel consumption and delivery delays by autonomously adapting to real-time disruptions. In this workflow, a **logistics control tower orchestrator** monitors the entire active fleet in real time, while a **disruption sensing agent** scans external data feeds for traffic accidents, severe weather, or port congestion. When a delay risk is identified, a **route planning agent** runs thousands of simulations to find the optimal alternative path that balances fuel costs, driver Hours of Service (HOS), and delivery windows. The **orchestrator** automatically pushes the new route to the driver’s navigation system and triggers a **customer communication agent** to update the consignee with a revised ETA.



The benefits are tangible. By compressing response cycles from hours to minutes, this approach can reduce fuel consumption by 15%, improves on-time performance, and enhances network resilience—demonstrating how orchestration at machine speed transforms disruption management from a reactive process into a core operational capability. This example operationalizes the Supply Chain Orchestrator found in the Agentic Genome Map, ensuring resilience against the supply chain shocks and highlighting dynamic route optimization as a key capability.

Autonomous Predictive Fleet Maintenance

In an agentic transportation enterprise, fleet maintenance shifts from a reactive, schedule-driven activity to an autonomous, condition-based operating model. By moving to Autonomous Predictive Maintenance, you're essentially turning a fleet of vehicles into a self-healing network. This shift eliminates unplanned downtime and extends asset lifecycles through self-executing, real-time interventions.

Autonomous Predictive Fleet Maintenance Workflow

Component	Role	Intelligence Layer
Maintenance Orchestrator	Mission Control	Synchronizes the vehicle's telemetry from IoT sensors (engine temperature, vibration, tire pressure), the driver's hours of service (HOS), and the repair facility's capacity into a single automated event.
Diagnostic Agent	Failure Prediction	Analyzes high-frequency vibration and thermal telemetry to identify "fingerprints" of failure, calculating probability windows (e.g., 90% chance of failure within 48 hours).
Scheduling Agent	Logistics Optimization	Accesses real-time GPS and telematics to find the path of least resistance—booking service at a shop already on the current route to avoid route disruptions.
Procurement Agent	Just-in-Time Logistics	Interfaces with parts inventory systems to pull the specific component to the service center, ensuring zero-wait time upon vehicle arrival.

Fleet Operational Wins:

- **Zero "Surprise" Downtime:** You stop fixing things because they *broke* and start fixing them because the data *knew* they would, keeping the fleet on the road.
- **Minimized Disruption:** Instead of towing a broken truck from the side of a highway, the truck checks in during a scheduled rest period or near its destination.
- **Extended Asset Life:** By addressing anomalies (like high engine heat) before they cause catastrophic damage, the residual value of the fleet stays higher for longer.

A closed-loop maintenance workflow can reduce unplanned breakdowns by 30-50%, extend asset life by 20-40%, and shift maintenance from a cost center to a strategic value driver.

Agentic Freight Documentation and Customs Clearance

This workflow addresses one of the most persistent bottlenecks in global trade: the “paper wall.” By replacing manual data entry with an agentic sequence, a shipment can be cleared by customs before the vessel even docks, virtually eliminating costly port storage fees.

Agentic Freight Documentation and Customs Clearance Workflow

Component	Role	Intelligence Layer
Trade Orchestrator	Lifecycle Coordination	Initiates the documentation process as soon as a shipment is booked. Synchronizes the documentation timeline with the physical movement of freight, ensuring no step is missed from booking to border crossing.
Document Classification Agent	Data Extraction & Structuring	Employs Advanced OCR and NLP to transform unstructured PDFs (commercial invoices) and photos into clean, categorized data sets (e.g., matching a line item to a specific HS [Harmonized System] Code).
Compliance Agent	Regulatory Validation	Cross-references extracted data against a live database of International Trade Regulations, sanctions, and tariffs to identify legal or financial risks.
Customs Filing Agent	Submission & Recording	Interfaces directly with government Electronic Data Interchange (EDI) systems to file declarations and secures the record in an audit-ready log.


The Frictionless Border Advantage:

- **Demurrage Elimination:** The value of this agentic AI application is potential accelerated customs clearance by up to 60%, significantly reducing the risk of goods being held at ports and incurring costly demurrage fees.
- **Audit-Ready Accuracy:** Since agents follow a standardized rule set and log every action, the risk of human error in HS Code classification—which can lead to massive fines—is significantly reduced.
- **Global Scalability:** Sovereign deployment models ensure that sensitive trade data remains governed within jurisdictional boundaries. The system can swap out “Regulation Modules” based on the destination country, allowing a small logistics team to handle trade across dozens of different jurisdictions effortlessly.

When the agentic genome is applied across transportation and logistic operations, the system shifts from reactive to anticipatory behavior. The result is not just faster throughput. It is a more resilient system—one that learns from each disruption and adapts its workflows over time, as illustrated in the following case study.



An Australian Port Authority



I have confidence knowing everything has been captured.

– Chief Information Officer, an Australian Port Authority

An Australian Port Authority manages some of the nation's busiest maritime gateways, coordinating berth scheduling, cargo handling, compliance, and maintenance across a complex ecosystem of ships, terminals, and intermodal networks. To modernize core operations and support rapid decision-making, the authority adopted an integrated content and process platform that consolidated incident management, asset records, and operational workflows.

Building on this system, the port authority is moving toward an agentic operating model that accelerates responsiveness and resilience. Intelligent sensing agents continuously ingest real-time telemetry from equipment sensors, vessel arrival feeds, and environmental data, correlating anomalies with historical failure patterns stored in governed content systems. An orchestration layer coordinates diagnostic, scheduling, and maintenance agents to proactively initiate condition-based interventions, while dynamic planning agents reroute workflows and resources to minimize disruption. When regulatory compliance flags arise—such as cargo manifest inconsistencies or safety deviations—compliance agents surface issues to human stewards who retain oversight and authority. This orchestration compresses cycles of detection, planning, and action without sacrificing governance or auditability.

The outcome is a port that is not just digitized but anticipatory. Unplanned downtime is reduced, assets are maintained proactively, berth and crane utilization become more fluid, and disruptions to cargo flows are managed with agility. By aligning agentic intelligence with a trusted content foundation and human-in-command governance, a major Australian port authority demonstrates how regulated infrastructure operators can transform from reactive service organizations into coordinated, resilient systems that deliver both operational excellence and strategic value in an era defined by speed, complexity, and volatility.

Automotive

The automotive industry operates at the intersection of complex global supply chains, highly regulated manufacturing environments, rapidly evolving vehicle software, and increasingly data-driven customer experiences. OEMs and suppliers must coordinate design, production, quality, logistics, recalls, and compliance across thousands of partners, while responding in near real time to disruptions ranging from parts shortages and factory downtime to regulatory changes and shifting consumer demand. Yet data remains fragmented across PLM, ERP, MES, quality, and supplier networks, and many workflows still depend on brittle handoffs and delayed human intervention.

Applying agentic AI through the Agentic Genome embeds reasoning and execution directly into automotive operating models. By grounding agents in trusted enterprise data and coordinating them through orchestrators across engineering, manufacturing, supply chain, and aftersales workflows, automotive organizations can enable end-to-end processes that adapt dynamically, improve quality and safety, accelerate throughput, and increase resilience across the vehicle lifecycle.

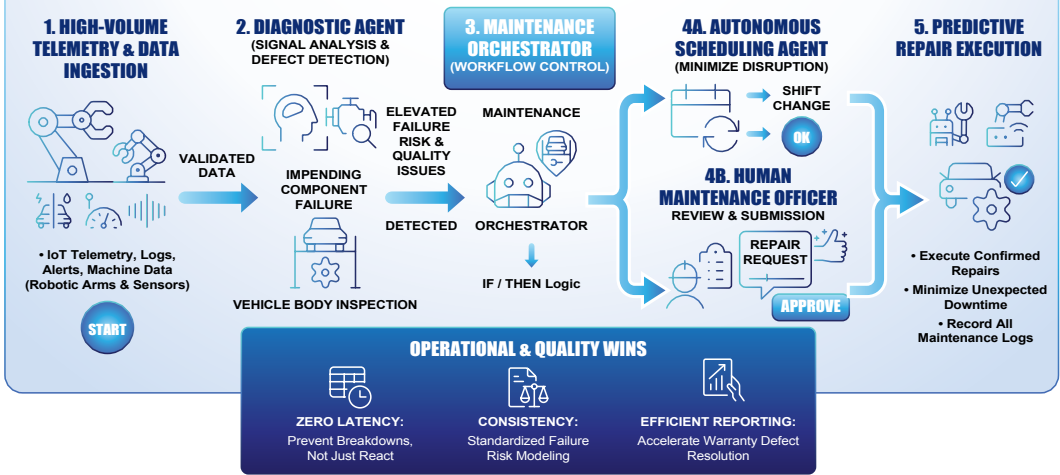
Predictive Manufacturing Asset Maintenance and Quality Assurance

To minimize production downtime and warranty defects, manufacturers can apply predictive manufacturing asset maintenance and quality assurance through agentic AI workflows that anticipate failures before they occur.

In this workflow, a **maintenance orchestrator** ingests high-volume telemetry and machine-generated data, including logs and alerts, from robotic arms and assembly line sensors. A **diagnostic agent** analyzes these signals for vibration anomalies or heat signatures that indicate impending component failure, while **computer vision agents** simultaneously inspect vehicle bodies on the production line for defects. When elevated failure risk is detected, a **scheduling agent** autonomously books a maintenance slot during a shift change to avoid disrupting production.

This model shifts maintenance from a fixed schedule to a condition-based approach, extending component life, preventing unplanned downtime, and ensuring consistently high-quality production outputs. It reflects practices seen in industrial environments that use wireless-enabled sensors and IoT models to predict repair needs, and it aligns with the operational technology layer of sovereign cloud architectures that bridge infrastructure and applications.

AGENTIC WORKFLOW: PREDICTIVE MANUFACTURING ASSET MAINTENANCE & QUALITY ASSURANCE



Intelligent Aftermarket Parts and Logistics Optimization

To optimize aftermarket parts availability and reduce inventory carrying costs, automotive organizations can deploy agentic AI workflows that dynamically predict demand across the dealer network.

Intelligent Aftermarket Parts and Logistics Workflow

Component	Role	Intelligence Layer
Logistics Orchestrator	Dynamic Network Control	Integrates data from back-end ERP systems with front-end sales data from dealers. Orchestrates the end-to-end flow, deciding whether to move existing stock across the network or order new supply to minimize total landed cost.
Demand Sensing Agent	Pattern and Trend Analysis	Correlates front-end dealer sales with historical seasonality and live market trends to detect shifts in part consumption before they become backorders.
Replenishment Agent	Autonomous Inventory Balancing	Executes Geo-Spatial Optimization to pull stock from the closest hub or triggers automated vendor POs based on pre-set economic order quantities (EOQ).

Strategic Value for Automotive Ops:

- **Minimizing “Lost Sales”:** Nothing kills customer loyalty like a “part backordered” notification. This system ensures the part is already in transit before the dealer even realizes they’re low.

- **Inventory Velocity:** By accurately predicting which parts are “movers,” you can reduce the footprint of slow-moving “dust-gatherers” in regional centers, freeing up working capital.
- **Proactive Supplier Relations:** Suppliers receive more accurate, earlier signals, allowing them to optimize their own production and potentially offering better terms for the predictable volume.

This workflow transforms the aftermarket supply chain from a “guess and check” model into a precision-guided replenishment engine. By bridging the gap between dealer-level sales and regional distribution, the system ensures that high-demand parts are always within reach without bloating the balance sheet with excess inventory.

Agentic Warranty Claims Processing and Fraud Detection

To reduce warranty leakage and fraudulent claims, organizations can apply agentic warranty claims processing and fraud detection through autonomous auditing of service data.

Agentic Warranty Claims and Fraud Detection Workflow

Component	Role	Intelligence Layer
Warranty Orchestrator	Claim Ingestion & Lifecycle	Acts as the central hub, using NLP to parse messy, unstructured mechanic notes and link them to standardized parts invoices.
Pattern Recognition Agent	Fraud & Anomaly Detection	Employs predictive analytics to spot red flags like mileage "jumping," unusual repair frequencies for specific VINs, or clusters of identical claims from a single dealer.
Compliance Agent	Policy Verification	A rule-based reasoning engine that checks the specific claim against the vehicle's exact warranty coverage, expiration dates, and previous repair history.
Transaction Agent	Automated Settlement	Triggers the financial execution for valid, low-risk claims, ensuring dealers are paid faster while flagging only the truly suspicious outliers.

This agentic AI workflow automates the reconciliation of financial records and claims, reducing manual labor and preventing revenue loss from invalid or fraudulent claims, while providing early visibility into emerging issues to pinpoint responsible business units. It adapts proven uses of AI for fraud detection and streamlined fleet administration and aligns with the Finance and Audit Agents model in the Agentic Genome Map, particularly the role of a dedicated Fraud Prevention Agent.

Just as a pattern recognition agent identifies anomalies in repair frequencies to stop warranty leakage, an identity orchestrator monitors access patterns to prevent unauthorized entries. By examining the digital transformation of this Automobile Club (below), we see how a centralized identity layer acts as a ‘security sensor,’ allowing specialized agents to grant seamless access to members while instantly flagging the anomalous behaviors that signal a breach.



An Automobile Club



A major mobility and membership services organization serving more than 13.5 million members needed to unify and secure customer interactions across a wide range of digital and physical channels—including mobile apps, call centers, branch visits, and web portals—while reducing complexity, improving security, and delivering seamless experiences. Existing identity and access systems were disjointed, slowing service delivery and creating friction for members trying to access multiple products and support services from different business units and devices.

The organization deployed a unified digital identity and access platform across its enterprise, consolidating member profiles and authentication into a single, secure identity service. In an agentic architecture, an identity orchestrator would serve as the central controller for identity workflows, while member context agents continuously ingest and normalize interaction data across channels. Complementary compliance and access agents would enforce policy in real time—granting appropriate access, detecting anomalous access patterns, and triggering alerts or escalations when needed—ensuring secure and consistent member engagement across all touchpoints.

By centralizing identity management and streamlining access paths, the organization significantly reduced digital complexity for its members. Members could securely access any service from any authenticated channel or device with a single digital identity, improving the user experience and reducing the operational overhead associated with fragmented sign-on and access policies. The platform establishes a trusted foundation for future digital services, including AI-driven automation of customer requests and personalized member engagements.

Healthcare

Healthcare organizations operate at the intersection of complex clinical workflows, fragmented information systems, and stringent regulatory obligations, where timely decisions depend on accurate, complete, and trusted data. Care delivery spans hospitals, clinics, labs, imaging centers, payers, and public health agencies, yet patient records, orders, referrals, and care plans are often distributed across siloed Electronic Health Records (EHRs), departmental systems, and unstructured content such as notes, images, and reports. This fragmentation slows coordination of care, increases administrative burden, and raises the risk of errors.

At the same time, healthcare operates under some of the most demanding governance and compliance requirements of any sector. Privacy laws, consent management, data residency rules, and strict auditability standards dictate how sensitive patient information may be accessed, shared, and acted upon. New regulatory frameworks—such as the European Union’s AI Act, which classifies many clinical AI systems as high-risk—extend these obligations to intelligent systems themselves, requiring transparency, human oversight, traceability, and lifecycle governance for AI deployed in medical decision-making and healthcare operations.

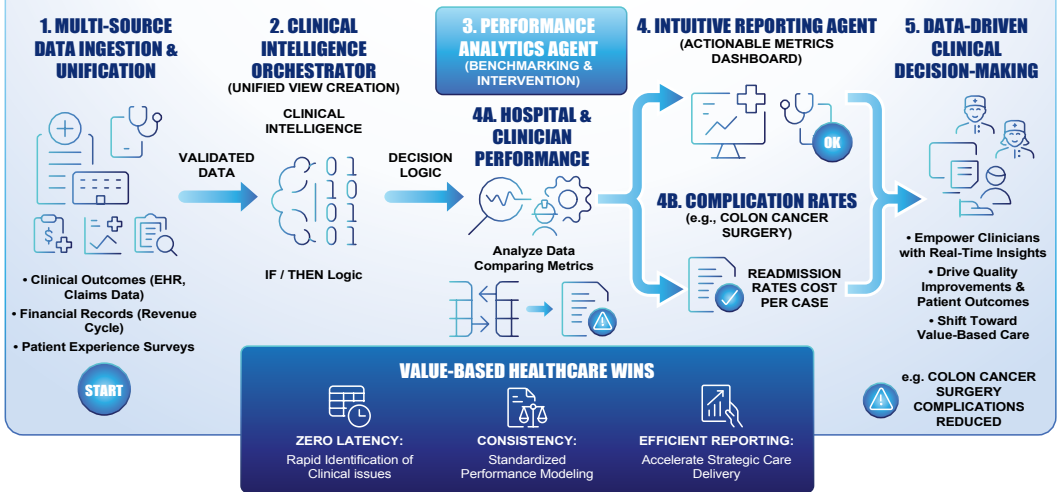
In this environment, the value of agentic AI depends not only on its ability to automate clinical and administrative workflows, but on the governance architecture surrounding it. Agentic systems must operate on trusted health information, respect patient consent and regulatory boundaries, and record every recommendation, action, and data access within a verifiable audit trail. Only when intelligent automation is paired with rigorous governance can healthcare organizations safely scale AI-enabled care coordination, clinical support, and operational efficiency while maintaining regulatory compliance and patient trust.

As illustrated in the examples below, applying an agentic AI workflow offers a path to orchestrate these data and processes across clinical, operational, and compliance domains, embedding governed reasoning and execution into care delivery while preserving privacy, safety, and accountability.

Value-Based Care Performance Analytics

To improve clinical outcomes and operational efficiency through benchmarking across the care continuum, value-based care performance analytics can be implemented as an agentic, data-driven workflow.

AGENTIC AI WORKFLOW: VALUE-BASED CARE PERFORMANCE ANALYTICS



In this agentic workflow a **clinical intelligence orchestrator** aggregates data from disparate sources, including clinical outcomes, financial records, and patient experience surveys, creating a unified view of performance. An **analytics agent** analyzes this information to compare outcomes across hospitals and clinicians, identifying best practices as well as areas requiring intervention. A **reporting agent** generates intuitive dashboards for healthcare professionals that surface actionable metrics such as complication rates for specific procedures.

This approach facilitates the shift toward value-based healthcare by enabling data-driven decision-making for clinicians and patients and has been shown to reduce complications substantially for colon cancer surgeries, illustrating how agentic AI bridges the gap between raw data and actionable clinical insight across the data and AI value chain.

Predictive Medical Imaging Asset Maintenance

To detect equipment failures before they impact patient care and to optimize service delivery, predictive medical imaging asset maintenance can be implemented as an agentic, continuously operating workflow.

Predictive Medical Imaging Asset Maintenance Workflow

Component	Role	Intelligence Layer
Maintenance Orchestrator	High-Volume Data Ingestion	Acts as the Central Nervous System, processing 1M+ events and 200k sensor readings daily per MRI unit to maintain a real-time state of the machine. Coordinates downstream actions.
Diagnostic Agent	Failure Pattern Recognition	Uses anomaly detection algorithms to scan telemetry for “silent” indicators—like subtle power fluctuations or cooling gradients—that precede a hard failure.
Service Agent	Autonomous Resolution & Dispatch	Employs prescriptive analytics to decide if a remote software patch can fix the issue or if a physical technician must be booked for an off-peak window.
Digital Twin (Edge)	Virtual Modeling	Maintains a synchronized digital replica of the specific MRI unit to simulate how current stress levels will impact the Mean Time To Failure (MTTF).

Healthcare Operational Wins:

- **Patient Continuity:** By predicting a “magnet quench” or cooling failure days in advance, the hospital can move maintenance to a Sunday night, ensuring no patient’s scan is canceled.
- **First-Time Fix Rate:** The diagnostic agent tells the technician exactly which part is failing before they leave the warehouse, ensuring they arrive with the right tools.
- **Remote Recovery:** Approximately 30-40% of modern imaging faults can be resolved via software resets or recalibrations initiated autonomously by the Service Agent.

In the Philips Healthcare case study in Chapter 6, we saw how this approach reduced equipment downtime by 30% and achieved an 84% first-time fix rate for onsite issues, directly enhancing patient outcomes and patient satisfaction, while demonstrating the shift from reactive support models to autonomous operations.

Agentic Patient Triage and Benefits Navigation

To enhance patient access to care and streamline administrative triage through intelligent automation, agentic patient triage and benefits navigation can be implemented as an orchestrated workflow across clinical and administrative systems.

Agentic Patient Triage and Benefits Navigation Workflow

Component	Role	Intelligence Layer
Patient Services Orchestrator	Channel & Workflow Management	Acts as the "Traffic Controller," managing secure portal data and ensuring seamless handoffs between clinical and administrative agents.
Triage Agent	Intent Classification	Employs Natural Language Understanding (NLU) to distinguish between urgent clinical symptoms, routine scheduling, and complex insurance questions.
Benefits Navigation Agent	Eligibility & Financial Clearance	Interfaces with Payer API Gateways to autonomously verify coverage, calculate co-pays, and secure prior authorizations in real-time.
EHR Scheduling Agent	Resource Optimization	Directly queries the Electronic Health Record (EHR) calendar to book slots that match the patient's urgency and the provider's specialty.

Key Healthcare Outcomes:

- **Reduced Time-to-Care:** By automating eligibility checks and scheduling, patients can often be seen days or weeks sooner than through a traditional manual call center.
- **Clinical Prioritization:** The NLU layer can flag high-risk symptoms (e.g., chest pain vs. a routine check-up) to ensure clinical resources are immediately alerted.
- **Administrative Efficiency:** Automating insurance verification reduces the rate of claim denials and lightens the load on front-desk staff, allowing them to focus on in-person patient hospitality.

This workflow reflects the agentic use cases described in sovereign EAI architectures in Chapter 3 and the principles demonstrated in large-scale request classification and automated task routing.

Whether classifying a patient's intent through NLU or identifying software defects through advanced machine learning, the core requirement for sovereign healthcare AI is absolute precision. By adopting AI-based functional testing, Roche Diagnostics (below) has replaced brittle, manual intervention with a consistent, automated control plane that mirrors the reliability required for end-to-end patient navigation.

Roche Diagnostics Shanghai



The AI-driven Functional Testing capabilities have drastically reduced test creation time and test maintenance work, while improving test reuse. This enhances our test coverage and increases our test asset resilience.

– Digital Solutions Manager of Commercial Innovation Department, Roche Diagnostics

Roche Diagnostics develops and delivers innovative, cost-effective, timely, and reliable diagnostic systems and solutions to support early detection and prevention of diseases. Its 2,800 employees aim to improve people's quality of life and reduce social medical costs. As an organization, Roche Diagnostics has embraced an agile development method to meet the increased digitization of many enterprises. Its Research and Development (R&D) team was given the challenge to deliver new software releases faster and more efficiently, while complying with the very specific quality requirements of the medical industry.

They selected a Functional Testing solution with advanced object recognition technology, enabled through AI-based machine learning and advanced OCR, which is proving very useful. By integrating Functional Testing with mature automation and development tools, users can perform fully automated tests without any manual intervention. This ensures testing efficiency and consistency.

The AI-driven capabilities have drastically reduced test creation time and test maintenance work, while improving test reuse. The introduction of Functional Testing has reduced regression test times by 90% through automation and the effective reuse of test cases. Whereas manual testing would take up to 12 hours for some applications, this is now done in a little over one hour. Test automation guarantees the core testing process and allows the development and testing teams to focus their work on adding true value to the business, through product innovation. This supports better job satisfaction and a more efficient use of valuable skills within the organization.

Pharmaceutical

Pharmaceutical organizations operate in one of the most highly regulated and documentation-intensive industries in the world, where every stage of the value chain—from research and clinical trials to manufacturing, safety reporting, and regulatory submission—must be supported by rigorous evidence, audit trails, and compliance with evolving regulatory frameworks.

Clinical development alone generates vast volumes of semi-structured documentation spanning protocols, consent forms, investigator files, adverse event reports, and regulatory submissions across a complex ecosystem of research partners, contract research organizations (CROs), clinical sites, regulators, and manufacturers. As organizations explore agentic AI to accelerate research and operational workflows, these systems must operate within strict governance frameworks—ensuring transparency, human oversight, and full traceability of automated decisions in accordance with global regulatory requirements, including emerging controls such as the European Union’s AI Act governing high-risk AI systems in health and safety domains.

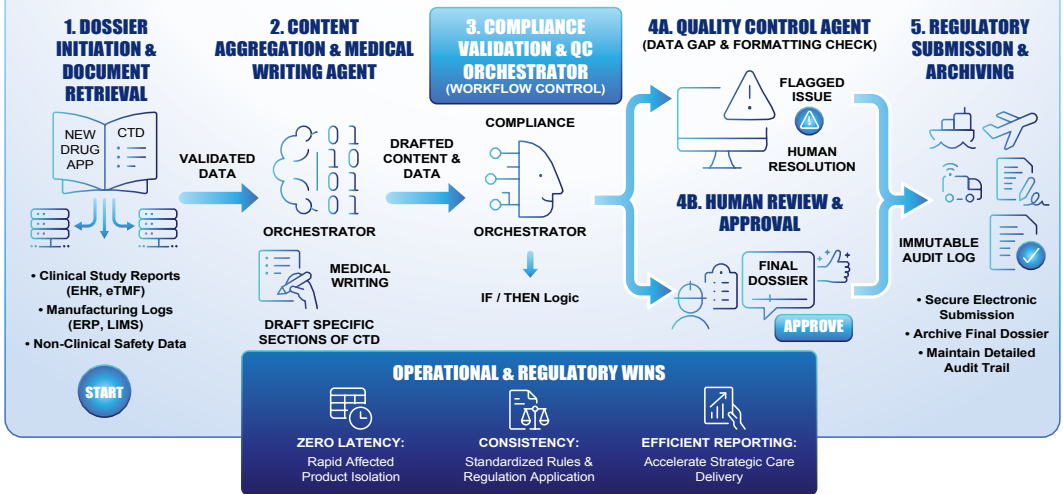
Applying the Agentic AI Genome map provides a structured way to orchestrate intelligence across this ecosystem—embedding agents and orchestrators into clinical, regulatory, safety, and manufacturing workflows so that reasoning and execution can be governed, auditable, and scaled across the pharmaceutical operating model.

Autonomous Regulatory Affairs and Submission Management

Automated regulatory affairs and submission management strives to automate the compilation and validation of regulatory dossiers in order to minimize submission rejection risks for new drug applications.

In this workflow, a **regulatory orchestrator** initiates the submission process for a new drug application. A **content aggregation agent** retrieves required documents from disparate silos, including clinical study reports, manufacturing logs, and non-clinical safety data. A **medical writing agent** drafts specific sections of the Common Technical Document (CTD), while a **compliance validation agent** cross-references the content against the latest health authority requirements across different regions. Any data gaps or formatting inconsistencies are flagged by a **quality control agent** for human resolution before the final dossier is generated.

AGENTIC WORKFLOW: AUTONOMOUS REGULATORY AFFAIRS & SUBMISSION MANAGEMENT



This agentic approach streamlines a complex, highly regulated submission process by ensuring accuracy and compliance across multiple jurisdictions, accelerating approval timelines, and supporting faster revenue generation. It leverages the Governance Orchestrator to oversee compliance as a continuous, orchestrated capability rather than a one-time submission event.

Agentic Drug Discovery and Clinical Trial Optimization

This workflow represents the “Holy Grail” of the pharmaceutical industry: shortening the decade-long, multi-billion dollar journey from a lab bench to a patient’s bedside. By using agentic AI, researchers can move from manual hypothesis testing to high throughput *in silico* discovery.

Agentic Drug Discovery and Clinical Trial Optimization Workflow

Component	Role	Intelligence Layer
R&D Orchestrator	Pipeline Synchronization	Manages the secure flow of data between public knowledge bases (patents/literature) and Private Zones (compound libraries).
Literature Analysis Agent	Knowledge Synthesis	Utilizes LLMs to ingest millions of biomedical papers and patent databases, identifying hidden correlations and novel biomarkers.
Molecular Simulation Agent	Predictive Chemistry	Employs Physics-Informed Neural Networks (PINNs) in the secure Private Zone to model 3D protein-ligand interactions and predict binding affinity without physical synthesis.
Clinical Design Agent	Protocol Optimization	Analyzes historical trial failures and successes to refine inclusion/exclusion criteria, maximizing the probability of statistical significance.
Regulatory Agent	Continuous Compliance	A semantic reasoning engine that cross-references trial protocols against the latest FDA/EMA guidelines to ensure compliance-by-design.

Together, these coordinated agentic AI workflows reduce the drug discovery and planning phase from years to months, significantly lowering R&D costs and accelerating the delivery of life-saving therapies to patients. This agentic AI workflow leverages the Private Zone architecture described in the companion book in this series to protect high-value intellectual property such as molecular structures while enabling orchestrated intelligence across the R&D lifecycle.

Intelligent Batch Record Review and Release

The objective of intelligent batch record review and release is to ensure pharmaceutical product safety and expedite batch release by automating quality control and environmental monitoring.

Component	Role	Intelligence Layer
Quality Orchestrator	Real-Time Oversight	The “Master Controller” that synchronizes live production telemetry with laboratory results and documentation timelines and monitors the manufacturing process in real time.
IoT Agent	Telemetry Ingestion	Captures high-frequency data from floor sensors (temperature, pressure, mixing speeds) to ensure the physical environment stayed within validated states.
Document Analysis Agent	Data Digitization	Uses advanced OCR & NLP to transcribe handwritten operator signatures and notes from paper logs into structured, searchable digital data.
Anomaly Detection Agent	Golden Batch Comparison	A machine learning engine that compares current run data against the idealized historical “Golden Batch” to spot even micro-deviations in real time.
Release Agent	Documentation Synthesis	Auto-generates the Certificate of Analysis (CoA) and final release dossier, flagging “Release-Ready” status for the Qualified Person (QP).

Operational and Safety Wins:

- **Compressed Release Cycles:** Reduces the batch review cycles from weeks to hours by eliminating the manual “paper-shuffling” phase of quality assurance. Improves supply chain velocity and reducing the amount of working capital tied up in quarantined inventory.
- **Proactive Deviation Management:** Instead of discovering a temperature excursion three days after the batch is finished, the orchestrator flags it the second it happens, potentially saving the batch.
- **Audit-Proof Compliance:** Creates an immutable digital trail of every sensor reading and operator note, making regulatory audits (FDA/MHRA) significantly less labor-intensive.

In the highly regulated world of life sciences, agentic means more than just speed—it means immutable, audit-proof compliance across every lab and production floor. As seen in the transformation of this Major Pharma Company (below), establishing a centralized EIM platform would enable specialized governance agents to act as continuous sentinels, automatically tagging and routing regulated content to ensure that innovation velocity never comes at the cost of product safety.



A Major Pharmaceutical Company



A major pharmaceutical manufacturer producing more than 2,000 products, with sales in over 60 countries, faced a persistent operational bottleneck: paper-centric and siloed content, inconsistent governance, and manual coordination of highly regulated documents hindered quality assurance and slowed responsiveness to regulatory demands. With more than 45 million prescriptions filled annually in Canada alone, the company needed a unified information foundation to support quality, compliance, and fast access to authoritative documents across labs, production floors, and executive functions.

The organization deployed an EIM platform to centralize its most valuable business content into a strategic, governed repository. In an agentic architectural interpretation, a content orchestrator could continuously index and harmonize regulated documents—such as standard operating procedures (SOPs), test reports, and regulatory submissions—across systems and silos, while document classification agents would analyze incoming materials in real time to enforce version control, compliance policies, and audit trails. Governance agents would ensure global regulatory requirements are reflected in document lifecycles, automatically tagging, routing, and escalating content for review when deviations or gaps are detected. In this example, human experts remain in command of policy and exception resolution, while agents handle the heavy lifting of continuous compliance monitoring and document readiness.

By establishing a unified source of truth for regulated content, the company improved access to critical information across functions and accelerated decision-making. Centralized governance reduced the risk of deviations and ensured regulatory documents are accessible where they are needed—whether at a laboratory bench, the production line, or the boardroom—supporting both quality standards and innovation velocity.

Grounding regulated content in a governed, enterprise-wide system enabled a shift from reactive document retrieval and compliance support to proactive information orchestration. With a content platform capable of supporting agentic AI workflows—where agents continuously enforce policy, classify critical content, and surface actionable alerts—the organization is better equipped to accelerate time-to-market, uphold quality standards, and reduce manual compliance overhead. This positions the enterprise to evolve toward autonomous, orchestrated document management that scales with regulatory complexity and operational growth.

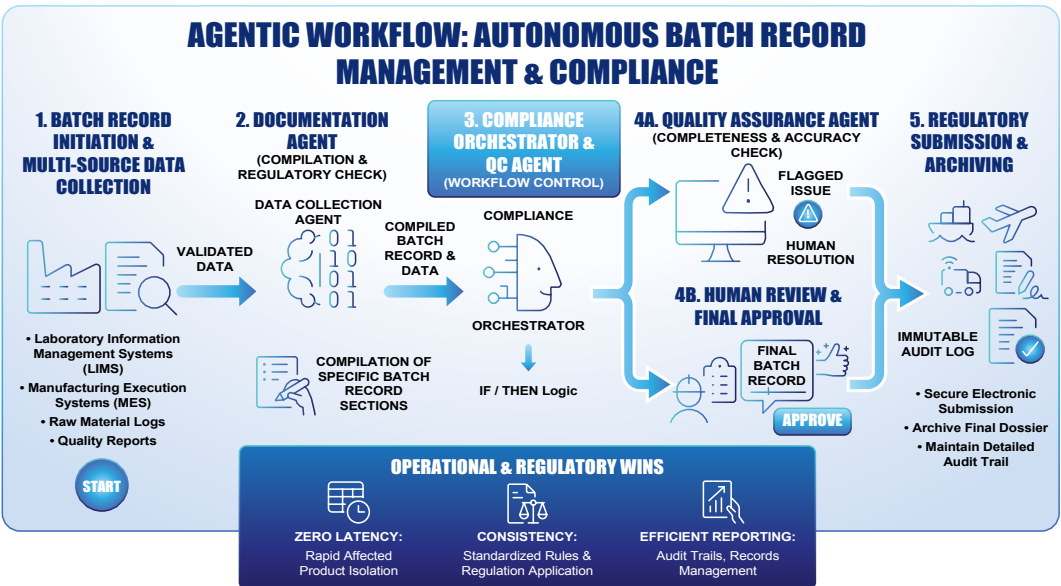
Process Manufacturing

Process manufacturing operates at the intersection of complex physical systems, tightly coupled supply chains, and stringent regulatory and quality requirements. Producers must manage continuous production flows, volatile commodity inputs, energy-intensive operations, and highly sensitive process parameters, all while meeting strict safety, environmental, and compliance standards. Data is abundant but fragmented across operational technology, laboratory systems, quality records, maintenance logs, and business networks, making it difficult to coordinate decisions across planning, production, quality, and compliance in real time. Disruptions propagate quickly, yet response remains constrained by manual handoffs, siloed systems, and reactive operating models that struggle to keep pace with variability at machine speed.

The Agentic Genome Map reframes process manufacturing around orchestrated intelligence rather than isolated automation. By grounding agents in trusted operational and quality data and coordinating them through orchestrators with human-in-command oversight, manufacturers can shift from point optimization to end-to-end, goal-directed workflows—embedding reasoning and execution directly into the operating model while preserving safety, compliance, and accountability.

Autonomous Batch Record Management and Compliance

Automating the documentation of batch records helps ensure regulatory compliance and streamlines audit readiness.



In this agentic workflow a **compliance orchestrator** initiates the documentation process for each production batch. Simultaneously, a **data collection agent** retrieves data from multiple systems, including laboratory information management systems (LIMS) and manufacturing execution systems (MES). A **documentation agent** then compiles this information into a comprehensive batch record, ensuring that all regulatory requirements are met. A **quality assurance agent** reviews the record for completeness and accuracy, flagging any discrepancies for human review before final approval.

By automating these workflows, organizations can reduce the time required for batch record review and approval from days to hours, accelerating product release and minimizing the risk of compliance violations. This approach utilizes the Governance Orchestrator concept to oversee compliance as described in Chapter 6 of the companion book in this series.

Agentic Process Control and Energy Optimization

This workflow represents the pinnacle of industrial autonomy, where the factory floor isn't just automated but it is also "aware." By balancing the competing priorities of energy costs, production yield, and worker safety in real-time, the system achieves a level of efficiency that manual oversight simply cannot match.

Intelligent Energy and Process Control Workflow

Component	Role	Intelligence Layer
Production Orchestrator	Multi-Stream Integration	Ingests high-frequency telemetry (temperature, pressure, flow) to create a real-time digital twin of the plant's operational state.
Energy Efficiency Agent	Cost-Load Balancing	Employs predictive economic modeling to cross-reference production needs with fluctuating energy grid prices, shifting heavy loads to "green" or cheap windows.
Process Control Agent	Dynamic Setpoint Tuning	A reinforcement learning engine that makes micro-adjustments to hardware (valves/heaters) to maintain the "sweet spot" for maximum yield.
Safety Agent	Critical Fail-Safe	A deterministic reasoning layer that bypasses all other agents to execute emergency shutdowns if physical thresholds (e.g., PSI limits) are breached.

Industrial “Triple-Bottom-Line” Wins:

- **Energy Arbitrage:** By slightly slowing non-critical processes when energy prices spike (and ramping up when they drop), the plant can reduce its utility bill by 15–20% without losing total output, while increasing production yield—to meet profitability and sustainability goals.
- **Waste Minimization:** The Process Control agent detects “drift” in heating elements long before a human operator would notice, preventing entire batches of raw material from being scorched or ruined.
- **Autonomous Safety:** Traditional “dumb” alarms wait for a disaster; this agentic safety layer can predict a breach based on the rate of change in pressure, acting seconds before a catastrophic failure occurs.

This workflow leverages the operational technology layer of sovereign cloud architectures that bridge infrastructure and applications.

Predictive Asset Maintenance and Reliability

This use case centers on minimizing unplanned downtime and maintenance costs by predicting equipment failures before they occur.

Predictive Asset Maintenance and Reliability Workflow

Component	Role	Intelligence Layer
Maintenance Orchestrator	Asset Health Monitoring	Acts as the central hub, continuously streaming multi-modal sensor data from pumps, compressors, and reactors into a unified health dashboard.
Diagnostic Agent	Signal & Anomaly Detection	Employs signal processing and pattern recognition to identify microscopic deviations in vibration, thermal, and acoustic signatures.
Reliability Agent	Prognostics & Life Prediction	Uses Remaining Useful Life (RUL) modeling to forecast exactly how many hours of operation remain before a component hits a critical failure threshold.
Work Order Agent	Autonomous Scheduling	Executes operational logic to cross-reference the maintenance need with the production schedule, generating work orders and booking repairs during “natural” downtime.



Operational and Financial Benefits:

- **Avoidance of “Secondary Damage”:** When a bearing fails, it often destroys the entire motor. By catching the bearing wear early, the repair cost is reduced by orders of magnitude.
- **Maintenance Cost Reduction:** Organizations typically see a 25–30% reduction in maintenance costs by eliminating unnecessary, schedule-based parts replacements.
- **Increased OEE (Overall Equipment Effectiveness):** By ensuring that machines only stop when planned, the plant’s total throughput and reliability increase significantly. Downtime can be reduced by up to 50%.

Strategic maintenance is built on the transition from reactive repairs to autonomous scheduling; a transition that is mirrored in the world of global logistics and partner onboarding. By consolidating 60 countries of fragmented trading data onto a single managed platform, this Steel and Mining Manufacturer (below) has moved beyond labor-intensive manual onboarding to an orchestrated ecosystem where agents can enforce compliance and handle exceptions at machine speed.



A Steel and Mining Company



A global steel and mining manufacturer operating across 60 countries faced growing complexity in onboarding new trading partners, managing a fragmented B2B network, and maintaining aging, in-house mainframe-based EDI applications. Partner onboarding was slow and labor-intensive, collaboration across suppliers and customers was inconsistent, and limited visibility across global transactions made it difficult to monitor performance and exceptions in real time. Running and maintaining these B2B applications in-house was costly, diverted scarce technical expertise, and constrained the company's ability to scale its digital trading network.

By moving to a managed B2B services model, the company established a unified digital backbone for partner connectivity and transaction exchange. On top of this foundation, agentic AI workflows can be introduced to orchestrate partner onboarding and transaction management end to end. Automated onboarding workflows reduce cycle times for bringing new partners onto the network and ensure consistent application of compliance and data quality rules across regions. Agentic monitoring and exception-handling workflows enable near real-time detection of transaction failures and partner connectivity issues, reducing disruption to supply chain and order fulfillment processes.

Entrusting Managed Services with the operation of more than 100,000 monthly EDI messages eliminated the need for in-house B2B technology expertise, reduced support and maintenance costs, and enabled faster onboarding of new trading partners. Core business processes became smoother and more reliable, and the organization gained increased visibility across its global business network. By consolidating its B2B trading networks onto a managed digital platform, the organization laid the groundwork for a more autonomous, scalable trading ecosystem. The result is a global business network that operates with greater transparency, reliability, and speed—supporting continuous expansion to additional trading partners while reducing operational burden and improving end-to-end visibility across international supply chains.

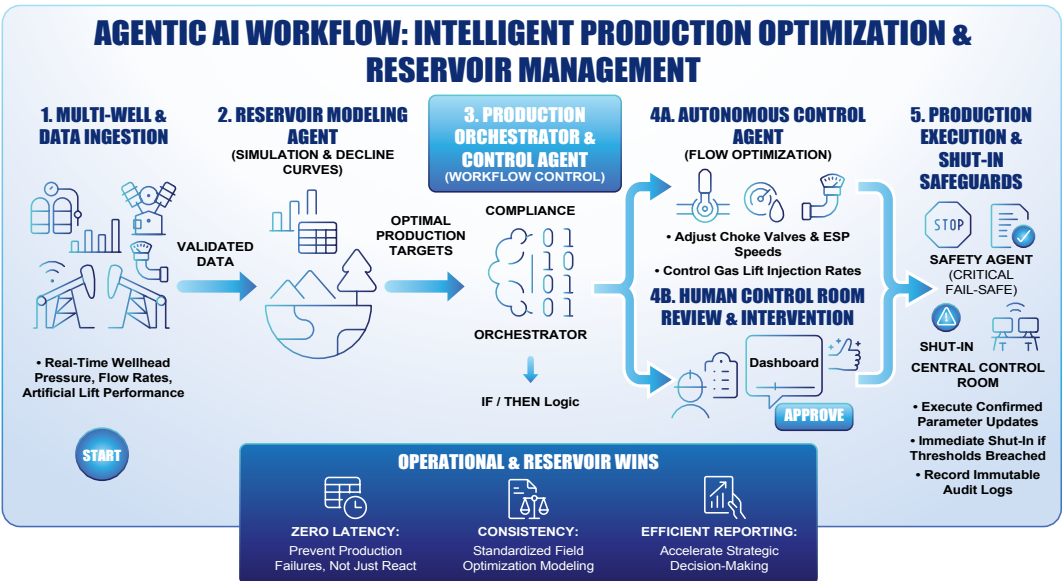
Oil and Gas

Oil and gas organizations operate some of the world’s most complex, capital-intensive projects, spanning exploration, drilling, production, refining, transportation, and decommissioning—often across multiple jurisdictions and in partnership with operators, contractors, regulators, and local stakeholders. These projects generate massive volumes of documents, drawings, permits, safety records, inspection reports, and contractual artifacts, much of it still managed through fragmented systems and paper-heavy processes that slow collaboration, obscure accountability, and increase compliance risk in a highly regulated environment. The outcome is operational friction at precisely the moments when speed, safety, and traceability matter most.

Applying the agentic genome map to oil and gas creates a path from disjointed automation to orchestrated intelligence. Agents grounded in governed project, asset, and regulatory data are coordinated by process and value-stream orchestrators to streamline permitting, document control, safety management, and supplier coordination—embedding reasoning and execution into the operating model while preserving regulatory compliance, auditability, and human-in-command oversight.

Intelligent Production Optimization and Reservoir Management

The objective of intelligent production optimization and reservoir management is to maximize barrel of oil equivalent output while adhering to geological and environmental constraints. In this agentic AI workflow a **production orchestrator** monitors real-time wellhead pressure, flow rates, and artificial lift performance across the field. A **reservoir modeling agent** continuously runs simulations to predict decline curves and identify optimal extraction rates. A **control agent** autonomously adjusts choke valves, gas lift injection rates, or electric submersible pump speeds to optimize flow without damaging the formation. If a safety threshold is breached, a **safety agent** triggers an immediate shut-in and alerts the central control room.



The value of this approach is increased production yield and optimized lifting costs, achieved through continuous micro-adjustments made 24/7 at a frequency and precision human operators cannot match, directly improving cash flow.

Autonomous Asset Integrity and Predictive Maintenance

To minimize nonproductive time and prevent catastrophic equipment failures in remote environments, agentic AI workflows coordinate predictive maintenance and asset integrity as a continuous, autonomous process.

Autonomous Asset Integrity and Predictive Maintenance Workflow

Component	Role	Intelligence Layer
Maintenance Orchestrator	SCADA & IoT Integration	Aggregates high-volume telemetry from remote assets such as offshore rigs, compressors, and turbines, serving as the digital command center that bridges the gap between field hardware and AI agents.
Diagnostic Agent	Multi-Modal Failure Detection	Analyzes complex vibration, thermal, and pressure "signatures" using pattern recognition to detect microscopic indicators of corrosion or mechanical fatigue.
Reliability Agent	Prognostic Modeling	Employs predictive probabilistic models to estimate the Remaining Useful Life (RUL) and the financial/safety risk of continued operation.
Logistics Agent	Supply Chain & Workforce Sync	A constraint-satisfaction engine that checks global inventory and technician certifications to execute a "pre-emptive strike" on the maintenance issue.

Critical Remote-Site Wins:

- **Elimination of Non-Productive Time (NPT):** In offshore environments, NPT can cost millions per day. This system ensures the rig keeps humming by replacing parts during scheduled logistical windows.
- **Autonomous Parts Logistics:** The system doesn't just say something is broken; it ensures the replacement part is already on a helicopter or supply ship before the failure window closes.
- **Enhanced Safety and Environmental Stewardship:** By detecting corrosion or mechanical wear early, you significantly reduce the risk of leaks, spills, or catastrophic structural failures in sensitive environments.

Automated Regulatory Compliance and HSE Monitoring

This workflow addresses the “Social License to Operate” by turning Health, Safety, and Environment (HSE) from a periodic manual audit into a real-time, autonomous oversight system. By bridging the gap between physical sensors and regulatory reporting, the enterprise moves from detecting violations to preventing them entirely.

Automated Regulatory Compliance and HSE Monitoring Workflow

Component	Role	Intelligence Layer
Compliance Orchestrator	Enterprise Data Aggregation	Acts as the central “Compliance Brain,” synchronizing emissions data, flaring volumes, and waste records across every facility into a single audit trail.
IoT Agent	Real-Time Environmental Sensing	Employs edge computing and anomaly detection to monitor methane intensity, gas leaks, and chemical spills from perimeter sensors at the source.
Documentation Agent	Regulatory Content Synthesis	Utilizes NLU (Natural Language Understanding) to map raw operational data into specific, localized legislative reporting formats (e.g., EPA or EEA standards).
Audit Agent	Policy Enforcement & Verification	A constraint-reasoning engine that compares real-time operations against specific permit conditions, identifying “Near-Misses” before they become violations.

Strategic HSE Advantages:

- **Fines and Liability Mitigation:** By identifying a methane leak or a flaring spike in minutes rather than during a monthly review, the organization avoids the massive financial penalties and license revocation associated with prolonged environmental breaches.
- **Audit Readiness (Always-On):** Instead of scrambling for data when a regulator arrives, the documentation agent ensures that a validated, accurate record is ready for submission at any second, reducing the administrative burden of compliance reporting.
- **Proactive Risk Management:** The audit agent identifies patterns of behavior that lead to violations, allowing the company to retrain staff or adjust equipment settings before a breach occurs.

To see how these theoretical frameworks of orchestrated intelligence translate into tangible operational gains, we can look to SOCAR Turkey, where the transition from fragmented manual processes to a unified, AI-driven service architecture is already redefining how the organization manages its vast scale and complex workforce.

SOCAR Turkey



The solution has increased our transparency and with its standardized processes we have improved control and governance. We can analyze and justify our workload and have a great platform for custom application development which has given us significant cost savings. We process up to 8,000 requests each month, with 96 percent submitted directly via the service portal or mobile app.

— IT Architecture Group Manager, SOCAR Turkey

Focusing on its main investment areas, namely Petro-Chemical products, refining, and natural gas trade and distribution, SOCAR Turkey provides support to enable Turkey to be an influential power across the international energy arena with its investments. As part of a corporate initiative, SOCAR consolidated their IT architecture to unify disparate business units, while modernizing key components. Their aim was to create a single service management platform serving all SOCAR Turkey companies and unifying IT processes across 5,500 end-users, 300+ service agents, over 100 functional groups and departments, and 7,000 IT assets. To do so, SOCAR selected an AI-powered, cloud-native software solution that integrates IT Service Management (ITSM), IT Asset Management (ITAM), and Enterprise Service Management (ESM).

The solution simplifies service management and automates workflows in a unified enterprise platform with a modern interface and mobile access. A single portal serves all needs: replacing faulty equipment, ID card renewal, mobile signature requests, and more. SOCAR uses it for time and capacity management within their enterprise application development teams, and it has been adopted by other non-IT units. Their insurance department uses it for insurance policy and claims tracking, their corporate fleet tracks and manages company car loans, their plant maintenance team uses the solution for inventory management, and their mobile management team tracks inventory and invoices. The Information Security team is evaluating assets in a custom app that supports the regulations they need to comply with from the Digital Transformation Office of the Republic of Turkey. The solution replaces stand-alone applications, which SOCAR would have had to purchase at significant cost to the organization.

The success of SOCAR's digital transformation underscores a pivotal shift in the Oil and Gas industry. By integrating IT service management with broader enterprise goals, the organization has laid the digital substrate required for the next generation of autonomous operations.

Energy and Utilities

Energy and utilities organizations operate some of the most complex, safety-critical infrastructure in the economy, spanning generation assets, transmission and distribution networks, substations, meters, and customer service systems. They must balance reliability, affordability, and decarbonization while operating under strict regulatory oversight, aging infrastructure, extreme weather events, and rising expectations for resilience and transparency. Data is generated across operational technology, field devices, enterprise systems, and partner networks, yet decisions are often slowed by fragmented systems, manual coordination across functions, and reactive operating models that struggle to keep pace with grid volatility and real-time demand.

Applying the Agentic Genome to energy and utilities provides a path to redesign these operations around orchestrated intelligence rather than isolated automation. By grounding agents in trusted grid, asset, and customer data; coordinating them through orchestration layers across generation, grid operations, and customer workflows; and governing execution with human-in-command controls, utilities can move toward predictive grid management, autonomous asset maintenance, real-time outage response, and compliant regulatory operations.

Autonomous Grid Optimization and Load Balancing

This workflow represents the transition to the smart grid, where the complexity of intermittent renewables (like wind and solar) and fluctuating demand requires sub-second decision-making. By moving from manual dispatch to an agentic model, the grid becomes a self-regulating ecosystem capable of preventing blackouts before they cascade.

Autonomous Grid Optimization and Load Balancing Workflow

Component	Role	Intelligence Layer
Grid Operations Orchestrator	Real-Time Network Ingestion	Acts as the "Digital Twin" of the entire utility network, syncing data from smart meters, substations, and distributed energy resources (DERs).
Forecasting Agent	Predictive Demand Modeling	Utilizes deep learning and meteorological integration to project load spikes and renewable generation dips based on hyper-local weather shifts.
Load Balancing Agent	Dynamic Supply Optimization	Employs constraint satisfaction algorithms to autonomously toggle battery storage, adjust flows, and initiate "Demand Response" to match supply.
Self-Healing Agent	Fault Isolation & Rerouting	A prescriptive AI layer that detects potential overloads or line failures and reconfigures the network topology to "airgap" issues and maintain service.

Strategic Grid Outcomes:

- **Renewable Integration:** Solves the “duck curve” problem by autonomously storing excess solar/wind energy in batteries and releasing it exactly when demand begins to peak. Optimizes energy flow without manual intervention while more effectively integrating intermittent renewable sources.
- **Resilience and Reliability:** The self-healing agent acts at machine speed, rerouting power in milliseconds to prevent a localized transformer failure from turning into a regional blackout.
- **Operational Efficiency:** Reduces the need for “Peaker Plants”—which are expensive and carbon-intensive—by using demand response and storage to smooth out the load curve.

This use case leverages the operational technology layer of sovereign cloud architectures that bridge infrastructure and applications.

Predictive Asset Maintenance for Utility Infrastructure

Predictive asset maintenance for utility infrastructure minimizes catastrophic asset failure and extends the life of critical infrastructure. By turning static infrastructure into a data-driven network, utilities can prevent catastrophic outages, such as transformer fires or turbine failures, while maximizing the lifespan of multi-million dollar assets.

Predictive Asset Maintenance for Utility Infrastructure Workflow

Component	Role	Intelligence Layer
Maintenance Orchestrator	Asset Health Oversight	Acts as the “Central Nervous System,” ingesting multi-modal telemetry from transformers, turbines, and lines to provide a unified fleet health view.
Diagnostic Agent	Multi-Sensory Analysis	Employs Computer Vision and Acoustic Signal Processing to detect microscopic thermal hotspots, abnormal vibration patterns, or “arcing” sounds.
Reliability Agent	Prognostics & Risk Scoring	A probabilistic reasoning engine that calculates Remaining Useful Life (RUL) and assigns a priority score based on the impact of a potential outage.
Field Service Agent	Autonomous Dispatch	Executes logistics and resource optimization to match the maintenance task with the nearest qualified crew, available parts, and optimal weather windows.

Strategic Utility Outcomes:

- **Avoidance of Secondary Damage:** By catching a bearing failure or a cooling leak in a turbine early, the utility avoids a catastrophic cascading failure that could destroy the entire unit.
- **Transition to Condition-Based Maintenance:** Moving away from arbitrary calendar-based schedules to actual health-based maintenance reduces maintenance waste and ensures parts are used for their full functional life.
- **Enhanced Grid Resilience:** By predicting failure points in the transmission network before they occur, the orchestrator helps maintain 99.99% uptime, even during peak load periods.

Automated Regulatory Compliance and Environmental Reporting

This agentic workflow automates the heavy lifting of ESG (Environmental, Social, and Governance) and regulatory mandates. By transforming compliance into a real-time oversight function, organizations can move beyond mere reporting to active environmental stewardship.

Automated Regulatory Compliance and Environmental Reporting Workflow

Component	Role	Intelligence Layer
Compliance Orchestrator	Multi-Facility Governance	Acts as the "Central Nervous System," synchronizing disparate data streams (emissions, water, waste) across all facilities to provide a holistic view of the organization's environmental footprint.
Data Collection Agent	Multi-Source Ingestion	Employs edge integration to ingest high-frequency data from IoT sensors and parse unstructured operational logs into structured formats.
Reporting Agent	Legislative Alignment	Utilizes Natural Language Generation (NLG) and semantic mapping to ensure all compiled data meets the specific technical standards of local and federal agencies.
Audit Agent	Continuous Verification	A rule-based reasoning engine that compares real-time metrics against permit limits and legal thresholds to identify "Near-Misses" before they become violations. Flags deviations for corrective action.



Core Operational Benefits:

- **Frictionless Reporting:** Reduces the “Reporting Tax” on operational staff by automating the manual gathering of data from spreadsheets and logs.
- **Proactive Mitigation:** The audit agent flags water usage or emissions spikes as they happen, allowing for corrective adjustments before a daily or monthly limit is breached.
- **Audit Defense:** Provides an immutable, time-stamped digital trail of all environmental data, significantly reducing the cost and stress of external government audits.

In a sovereign AI ecosystem, the transition from reactive reporting to proactive stewardship is built on the immediate availability of governed data. By examining the digital modernization at this Public Water and Gas Utility (below), we see how moving beyond paper-based archives allows specialized agents to prioritize emergency alerts and suggest workflow steps—mirroring the continuous verification logic used by audit agents to prevent environmental and operational failures.



A Metropolitan Utilities Company



A large public water and natural gas utility serving over 600,000 customers relied heavily on paper records and manual coordination to support field maintenance and emergency response. Field teams often spent hours on phone calls with back-office staff just to locate schematics, maintenance histories, and asset details stored in aging archives, making rapid response to infrastructure issues difficult and costly. Legacy paper schematics and siloed operational data hindered timely decision-making and increased the risk of delayed repairs and resource waste.

To modernize information access and streamline operations, the utility digitized nearly one million historical paper records and integrated them with core systems including ERP and geographic information data, creating a unified operational information backbone. On this foundation, an asset information orchestrator could be used to coordinate multiple agents to support field workflows: a document retrieval agent would provide engineers with instant, context-aware access to technical drawings and maintenance histories on mobile devices, a location intelligence agent would enrich records with GPS coordinates and spatial context, and an emergency response agent could then prioritize alerts and suggests workflow steps based on asset condition and incident severity. This would enable field crews to receive real-time, structured information at the point of need without back-office dependency.

With agentic AI workflows driving information delivery and execution, field engineers can make faster, data-driven decisions and locate assets with pinpoint accuracy—even under challenging conditions like snow-covered sites—leading to accelerated maintenance work and reduced internal coordination time. By transitioning from paper archives to digital, governed content, the utility would be able to reduce the operational costs of storing and managing physical records, saving an estimated \$300,000 annually while reclaiming valuable workspace. Cross-system orchestration ensures that critical maintenance decisions and workflows are supported by reliable data and consistent governance, enhancing operational resilience, minimizing downtime, and improving customer service delivery.

Public Sector – Federal

At the federal level, governments operate at massive scale across mission-critical programs that span benefits administration, taxation, border services, defense, public safety, and regulatory oversight. These organizations must coordinate thousands of agencies, legacy systems, and partner organizations while managing enormous volumes of sensitive records, case files, and operational data. As governments explore agentic AI to improve service delivery and operational efficiency, these systems must operate within strict governance frameworks shaped by regulations such as the EU AI Act, the U.S. Federal Information Security Management Act (FISMA) and Canada’s Directive on Automated Decision-Making, alongside national privacy and cybersecurity laws—ensuring transparency, human oversight, data sovereignty, and full auditability for every automated decision made within public-sector systems.

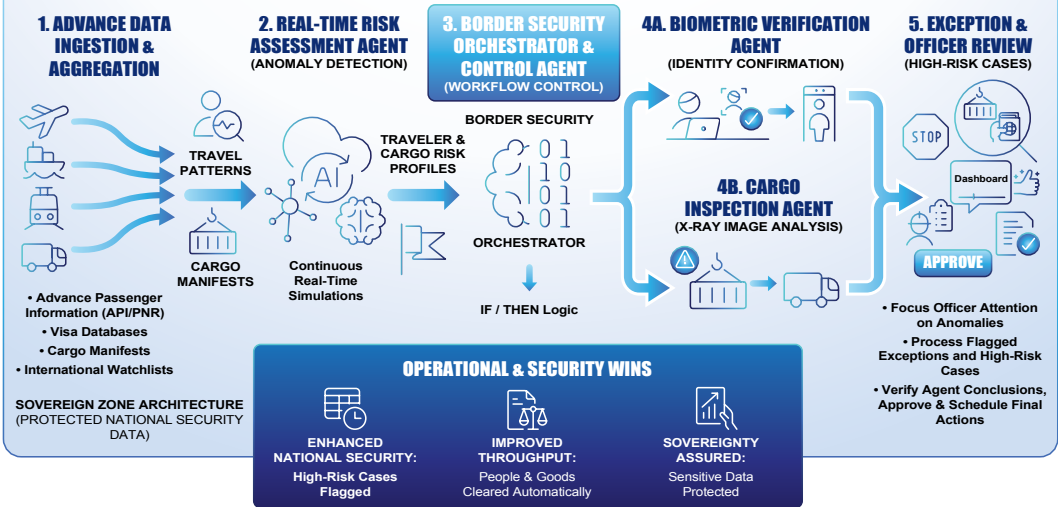
Applying the agentic AI genome to federal operations introduces a new operating model for content and process management. By coordinating domain-specific agents through process, mission, and enterprise orchestrators, federal agencies can streamline case handling, compliance enforcement, benefits delivery, and inter-agency coordination while preserving human-in-command oversight. In practice, agentic AI workflows enable faster service delivery, improved transparency, stronger compliance, and resilient operations across complex, multi-agency missions—embedding reasoning and execution directly into the machinery of government without compromising accountability, security, or public trust.

Autonomous Border Risk Assessment and Clearance

This goal of this agentic AI workflow is to enhance national security and border control by automating the screening of travelers and cargo.

A **border security orchestrator** integrates data from advance passenger information, visa databases, and watchlists to coordinate end-to-end risk assessment and clearance, while a **risk assessment agent** analyzes traveler patterns and cargo manifests in real time to identify high-risk anomalies. A **biometric verification agent** confirms identity at kiosks using facial recognition. A **cargo inspection agent** analyzes X-ray imagery for contraband.

AGENTIC AI WORKFLOW: AUTONOMOUS BORDER RISK ASSESSMENT & CLEARANCE



Low-risk travelers and goods are cleared automatically, allowing border officers to focus their attention on flagged exceptions and higher-risk cases. The effect is improved border throughput and security effectiveness, reducing wait times for legitimate travel and trade while intercepting threats more efficiently, with sensitive national security data protected through a Sovereign Zone architecture.

Agentic Benefits Eligibility and Case Management

This agentic workflow reimagines the public sector backlog by treating benefit applications as a high-velocity data stream rather than a manual paper trail. By automating the objective verification steps, human case workers are freed to focus exclusively on the nuanced, complex cases that require empathy and expert judgment.

Agentic Citizen Benefits Processing Workflow

Component	Role	Intelligence Layer
Benefits Orchestrator	Application Lifecycle Management	Manages the full application lifecycle for services such as employment insurance or disability support. Acts as the "Case Shepherd," ensuring data flows securely between intake, verification, and final decisioning while maintaining an immutable audit log.
Intake Agent	Data Extraction & Privacy Shielding	Uses Named Entity Recognition (NER) to pull facts from forms and automated redaction to mask Personally Identifiable Information (PII), ensuring privacy-by-design before data moves to downstream agents.

Component	Role	Intelligence Layer
Eligibility Agent	Cross-Agency Verification	Employs secure data interoperability to cross-reference application data against tax, residency, and citizenship databases in real time.
Case Management Agent	Decision Support & Routing	A logic-based reasoning engine that auto-approves "Perfect Match" cases and uses summarization LLMs to prepare condensed case packs for human review.

Public Sector Impact:

- **Radical Backlog Reduction:** By auto-approving straightforward claims (e.g., 60–70% of standard applications), the system clears the low-hanging fruit instantly, allowing humans to tackle the remaining backlog.
- **Privacy Compliance:** The intake agent’s ability to redact PII automatically ensures that sensitive data is only visible to authorized personnel, significantly reducing the risk of data breaches.
- **Consistency and Fairness:** Using a standardized eligibility agent ensures that every application is measured against the exact same criteria, eliminating human bias in the initial verification phase.

Intelligent Federal Procurement and Contract Management

This agentic AI workflow streamlines the procurement and acquisition processes to ensure fair competition for government contracts while reducing administrative burden.

Intelligent Federal Procurement and Contract Management Workflow

Component	Role	Intelligence Layer
Procurement Orchestrator	Lifecycle Governance	Acts as the "Master Controller," ensuring the procurement process adheres to the Federal Acquisition Regulation (FAR) and synchronizes data between agents.
Market Research Agent	Industry Intelligence	Employs predictive analytics and web scraping to analyze global pricing trends and identify qualified small or disadvantaged businesses.
Solicitation Agent	Document Generation	Uses GenAI (LLMs) to draft RFPs by cross-referencing successful historical templates with current technical requirements.
Evaluation Agent	Compliance & Scoring	A semantic analysis engine that objectively scores bid technicality and flags "Non-Compliant" submissions based on mandatory criteria.



Component	Role	Intelligence Layer
Contract Management Agent	Performance Oversight	Utilizes rules-based automation to track milestones and triggers payments only after verifying deliverables against the contract's Statement of Work (SOW).


Strategic Public Sector Wins:

- **Fairness and Transparency:** By using an objective evaluation agent, the human bias or perceived favoritism in scoring is significantly reduced, ensuring a level playing field for all vendors.
- **Administrative Velocity:** Automated RFP drafting and bid scoring can reduce the “Procurement-to-Award” cycle from months to weeks, getting critical infrastructure and services online faster.
- **Fiscal Accountability:** The contract management agent ensures that taxpayer funds are only released when milestones are truly met, preventing overpayments and identifying underperforming vendors in real-time.

While mission-specific agents handle borders and benefits, the internal resilience of the federal government depends on a robust digital backbone—a principle exemplified by the Court of Justice of the Federal District and Territories (TJDFT), which turned to AI and business process management capabilities to modernize its massive judicial infrastructure and reduce its administrative burden.

With over 7,600 civil servants, 378 judges, nearly 3,000 contractors, and 36 buildings, Court of Justice of the Federal District and Territories (TJDFT) relies on tools and IT support to keep its service running smoothly. The agency wanted to introduce more simplicity and more opportunity for end-users to self-service and collaborate with a modern service management solution. All issues were managed by 50 service agents in a 24/7 model that took too long to resolve. On average, less than nine percent of issues were resolved at first contact, resulting in loss of productivity.

Court of Justice of the Federal District and Territories



TJDFT opted for an AI-driven, cloud-native software suite that unifies IT Service Management (ITSM), IT Asset Management (ITAM), and Enterprise Service Management (ESM). The learning curve for the TJDFT team was very short, so they were up and running quickly, without a team of skilled developers or contractors. Within just a few months of introducing the solution, the first contact resolution rate went from 9 to over 45%; double the target set. Once the benefits of the new system became clear, many departments asked for non-IT processes to also be included, such as electronic voting at court committee meetings, stationery and office furniture order and supply, warehouse management, public servant travel on Court business, and to provide a repository of compliance documents, such as new laws and regulations.

During the pandemic, the team appreciated the solution's machine learning, smart search, and analytics capabilities. Users relied on the 'ask a friend' and 'feedback from other users' features to help each other within the user community. The AI-based virtual agent 'Max,' fully integrated into the knowledge base, supports user autonomy. When a question such as "how do I create a remote justice room?" is asked on the portal from a web client on a PC/laptop or a mobile device, Max leverages NLU (natural language understanding) and searches for documents created for this purpose. It uploads related articles dealing with the subject so that the user can select the most relevant with a click. The solution has empowered TJDFT users with a simplified self-service model, increased the first contact issue resolution rate, and achieved an ongoing cost saving of 40%.

The TJDFT experience highlights a vital tenet of the Public Sector Agentic Genome: simplicity is a prerequisite for scale. By moving beyond fragmented legacy tools to an integrated, AI-driven service model, the Court transitioned from reactive troubleshooting to proactive enablement. Ultimately, this transformation demonstrates that agentic workflows don't just solve IT tickets; they provide the institutional agility required to maintain public service continuity in an increasingly digital world.

Public Sector – State/Provincial

State and provincial governments sit on the front lines of public service delivery, administering benefits, licensing, healthcare programs, transportation, education, and regulatory oversight for millions of citizens and businesses. Their operations rely on high volumes of casework, fragmented legacy systems, and strict mandates around privacy, records retention, accessibility, and auditability. As these governments explore agentic AI to modernize services and reduce administrative burden, deployments must comply with governance frameworks such as U.S. state-level privacy regulations like the California Consumer Privacy Act (CCPA) and Canadian provincial statutes such as the Freedom of Information and Protection of Privacy Act (FIPPA)—ensuring transparency, explainability, human oversight, and full traceability for automated decisions that affect citizens and public services.

Applying the Agentic AI Genome at this level of government provides a path from disjointed automation to orchestrated intelligence, enabling agents and orchestrators to manage end-to-end workflows across departments, ground decisions in governed records and data, and enforce human-in-command oversight so services can be delivered faster, more accurately, and in compliance with regulatory and privacy obligations.

Agentic Court Case and Docket Management

The judicial system is often hampered by the sheer volume of paperwork and the logistical nightmare of manual scheduling. This workflow moves the court from a reactive, paper-driven model to a proactive agentic justice environment, significantly increasing the throughput of the legal process while ensuring data integrity.

Agentic Court Case & Docket Management

Component	Role	Intelligence Layer
Justice Orchestrator	Judicial Lifecycle Management	Acts as the central “Legal Brain,” managing the state transitions of a case from initial filing to final disposition.
Ingestion Agent	Document Digitization	Uses advanced OCR and computer vision to convert handwritten logs and physical evidence records into structured, searchable digital assets.
Verification Agent	Legal Data Validation	Employs entity resolution to cross-reference citations and case numbers against state criminal and traffic databases to prevent filing errors.
Scheduling Agent	Resource Optimization	Balances the calendars of judges, prosecutors, and public defenders with courtroom availability to optimize the docket.
Notification Agent	Stakeholder Synchronization	Utilizes real-time event triggers and automated messaging to ensure all parties are updated on docket changes via secure channels, keeping proceedings synchronized across the ecosystem.

This agentic AI workflow results in a significant reduction in manual data entry and scheduling conflicts, which shortens case lifecycles and frees court clerks to focus on complex legal processing rather than administrative work. This pattern mirrors outcomes in which AI-driven automation transforms courtroom operations and improves data integrity, demonstrated in the European Court of Human Rights case study in Chapter 5.

Autonomous Social Housing and Benefit Eligibility

This use case focuses on streamlining the processing of social housing and benefit applications to ensure timely support for vulnerable citizens through agentic AI workflows. By connecting directly to sovereign data sources, it ensures that critical support reaches vulnerable citizens with zero latency, while maintaining the highest standards of data privacy and statutory compliance.

Autonomous Social Housing and Benefit Eligibility Workflow

Component	Role	Intelligence Layer
Social Services Orchestrator	Application Lifecycle Governance	Acts as the “Central Command,” managing the state transitions of an application from initial intake to final disbursement or human escalation.
Intake Agent	Citizen Engagement & Data Capture	A conversational AI & OCR interface that guides citizens through complex forms and extracts structured data from uploaded identity or residency documents.
Eligibility Agent	Sovereign Data Validation	Utilizes secure API interoperability to cross-reference applicant data against tax, residency, and vital statistics databases within a protected sovereign zone.
Decision Agent	Statutory Rule Application	A deterministic reasoning engine that applies complex legal and policy frameworks to calculate benefit levels, flagging High-Risk or Edge Cases for human review.

Strategic Social Impact:

- **Equity through Velocity:** For citizens in precarious housing situations, the difference between weeks and days is critical. This workflow ensures that those who meet the criteria are approved almost instantly.
- **Fraud Prevention and Integrity:** By validating data against authoritative government sources (tax and vital statistics) rather than relying solely on self-reporting, the system significantly reduces the risk of fraudulent claims.
- **Human-in-the-Loop Optimization:** By automating 80–90% of standard applications, social workers can dedicate their expertise to the most complex and sensitive cases that require a high degree of empathy and nuanced judgment.

Intelligent Service Request Classification and Resolution

This agentic workflow transforms municipal service delivery from a “black hole” of manual requests into a high-velocity, transparent system. By automating the classification and tracking of citizen needs, local governments can ensure that urgent infrastructure issues—like a broken water main—are prioritized over routine inquiries without human intervention.

Intelligent Service Request Classification and Resolution Workflow

Component	Role	Intelligence Layer
Service Delivery Orchestrator	Lifecycle Management	Manages inbound citizen requests such as road maintenance, waste management, and permit inquiries. Acts as the “Citizen Experience Hub,” coordinating data flow between the public-facing portal and the backend departmental systems.
Classification Agent	Intent & Urgency Analysis	Employs Natural Language Processing (NLP) to categorize requests (e.g., “pothole” vs. “permit”) and assign a priority score based on public safety impact.
Routing Agent	Resource Allocation	Utilizes geospatial logic and system integration to dispatch requests to the appropriate field crew or department closest to the issue with the right tools.
Resolution Agent	Feedback Loop & Closing	A state-tracking engine that monitors the status of work orders, sends automated updates to the citizen, and uses sentiment analysis on feedback to improve service.

Strategic Municipal Wins

- **Reduced Response Latency:** Urgent requests are flagged and routed in seconds, bypassing intake bottleneck and allowing crews to be dispatched immediately. Can speed up request resolution by up to 60%.
- **Radical Transparency:** By providing real-time status updates, the city reduces the volume of follow-up calls to support centers, lowering administrative overhead. One-hundred percent visibility into performance metrics.
- **Data-Driven Urban Planning:** The system aggregates request patterns over time (e.g., recurring drainage issues in a specific ward), providing leadership with the insights needed for long-term infrastructure investment.

In an agentic ecosystem, the true value of a classification agent is its ability to bypass administrative bottlenecks by instantly identifying intent and urgency. For the following State Financial Regulator, the transition to an automated intake and validation agent mirrors this municipal model—transforming a slow, manual licensing process into a high-speed digital pipeline that scores risk and flags non-compliance in real time.



U.S. State Financial Regulator



A U.S. state financial regulator responsible for licensing and supervising traditional and emerging financial entities—including banks, fintechs, money transmitters, and digital insurers—faces operational complexity, fragmented systems, and high administrative costs. Multiple legacy platforms are typically used to process applications, which leads to slow turnaround times, costly maintenance, and a lack of developer cross-training and backup support when key personnel were unavailable.

To eliminate silos and streamline licensing workflows, the agency could deploy a unified agentic platform built on a low-code, enterprise content and process foundation as the platform for orchestration across agentic AI workflows. A licensing orchestrator would coordinate the end-to-end application lifecycle, invoking specialized agents that automate key steps, including: an intake agent to ingest and digitize application content; a validation agent to check completeness and compliance with regulatory criteria; a risk scoring agent to assess applicant risk profiles against internal and external datasets; and a decision assistance agent to prepare recommended outcomes, routing only ambiguous or high-risk cases to human adjudicators for review. Built on a single platform, these orchestrated agents eliminate redundant systems while preserving oversight and governance.

Key Operational Improvements:

- **Accelerated licensing:** Turnaround times for new and renewed licenses fall dramatically as structured, agentic workflows replace manual handoffs and disparate systems.
- **Consistent experience for users and staff:** A unified intake and decision workflow improves usability for both external applicants and internal regulators.
- **Cost consolidation:** Development and maintenance expenses drop as the agency consolidates multiple platforms into one governed system.
- **Developer agility and redundancy:** Teams trained on the shared platform ensure continuity of capability and reduce reliance on individual subject-matter experts.

By grounding intelligent agents in governed content and automating procedural steps through orchestrators, the regulator would achieve faster, more reliable licensing outcomes, reduced operational burden, and stronger overall service delivery without compromising regulatory integrity.

Public Sector – Local/Municipal

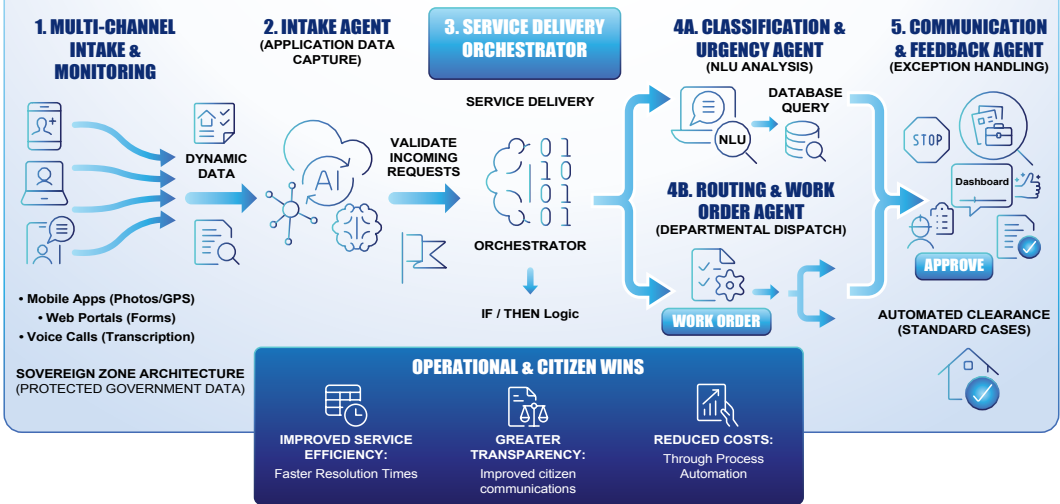
At the municipal and local level, governments sit closest to citizens—but often operate with the most constrained resources. City halls, utilities, planning departments, and service agencies manage high volumes of paper-heavy records, permits, inspections, complaints, and case files across fragmented legacy systems, while being expected to deliver fast, transparent service under strict privacy, records retention, and public accountability requirements. Work is highly cross-departmental, yet processes remain siloed, leading to backlogs, manual handoffs, limited visibility into case status, and uneven service quality across neighborhoods.

Applying the Agentic AI Genome to local government offers a path from fragmented digital services to orchestrated, end-to-end workflows. By grounding agents in governed municipal content, coordinating them through case and service orchestrators, and enforcing human-in-command oversight, municipalities can automate intake, triage, eligibility checks, scheduling, and compliance reporting across departments.

Intelligent Citizen Service Resolution 311

Service request management systems—commonly known as 311—respond to non-emergency requests from municipal citizens. Applying intelligent, autonomous processes to these systems improves both citizen satisfaction and operational efficiency.

AGENTIC AI WORKFLOW: INTELLIGENT SERVICE REQUEST MANAGEMENT (311)



In this workflow, a **service delivery orchestrator** ingests requests from mobile apps, web portals, and voice calls related to issues such as potholes, graffiti, or waste collection. A **classification agent** applies natural language understanding to determine the intent and urgency of each report, while a **routing agent** integrates with the city’s work order system to dispatch the appropriate department or field crew. Finally, a **communication agent** then autonomously updates residents on repair status and closes the ticket once work completion is verified.

The value of this workflow is a significant reduction in manual triage effort and faster resolution times, transforming citizen engagement from an opaque “black hole” into a transparent, responsive service experience.

Agentic Permitting and Zoning Compliance

This use case is introduced in Chapter 7 of the first book in this series under “Agentic Use Cases for Housing,” specifically citing application processing and permitting. The goal of agentic permitting and zoning compliance is to accelerate housing development and revenue generation by automating the building permit validation process.

Agentic Permitting and Zoning Compliance Workflow

Component	Role	Intelligence Layer
Permitting Orchestrator	Application Lifecycle Governance	Acts as the “Digital Concierge,” managing the secure intake of building applications and architectural BIM/ CAD files and coordinating hand-offs between validation agents.
Validation Agent	Technical Plan Review	Utilizes Computer Vision (CV) and spatial reasoning to cross-reference 2D/3D site plans against digitized zoning bylaws and structural building codes.
Compliance Agent	Credentials & Liability Verification	A data interoperability engine that queries provincial/ state licensing boards to confirm contractor standing, insurance coverage, and bonding in real time.
Approval & Payment Agent	Transactional Finalization	Employs deterministic logic to issue permits for standard applications and integrates with financial gateways to automate fee collection.

Strategic Housing and Economic Wins:

- **Accelerated Housing Starts:** By reducing the timeline from 180 days to 48 hours for standard builds, the city significantly lowers the carrying costs for developers, often translating to lower housing prices.
- **Consistency in Enforcement:** AI agents apply zoning bylaws with 100% consistency, eliminating the risk of human error or subjective interpretation that can lead to costly legal disputes or “variance creep.”
- **Resource Optimization:** Municipal planners are freed from checking setbacks on routine applications, allowing them to focus on high-impact projects like transit-oriented communities and heritage preservation.

Predictive Municipal Infrastructure Maintenance

This agentic workflow shifts municipal management from “fix-on-failure” to “fix-before-failure.” By treating city infrastructure as a living network of data, local governments can prevent the high social and financial costs of burst water mains or structural bridge fatigue.

Predictive Municipal Infrastructure Maintenance Workflow

Component	Role	Intelligence Layer
Infrastructure Orchestrator	Unified Asset Monitoring	Monitors real-time telemetry from geographically dispersed IoT sensors in pumps, bridges, and traffic signals.
Anomaly Detection Agent	Stress & Leak Identification	Employs time-series analysis and pattern recognition to distinguish between normal operational fluctuations and early indicators of structural or mechanical stress.
Prioritization Agent	Risk & Impact Assessment	A decision-support engine that calculates a “Criticality Score” based on the asset’s location, population density affected, and public safety risk.
Maintenance Agent	Logistics & Resource Dispatch	Executes operational logic to autonomously book inspection crews, procure specialized parts, and update the city’s capital budget forecast.

Strategic Public Infrastructure Wins:

- **Extending Asset Lifespan:** Small, proactive repairs (like sealing a minor leak or replacing a bearing) can add decades to the life of a multi-million dollar asset, maximizing the return on citizen tax dollars.
- **Budgetary Predictability:** Moving away from emergency repairs—which often cost 3–5x more than scheduled maintenance—allows for more stable and predictable municipal capital planning.
- **Public Safety & Continuity:** By catching structural stressors in bridges or traffic signal failures before they occur, the city eliminates the risk of accidents and the chaos of unplanned road or utility closures.

In the following case study about Auckland Transport, the integration of real-time video analytics and automated enforcement mirrors the core principles of the agentic genome—reducing operational costs and increasing public safety through the intelligent application of data at scale.

Auckland Transport



Our stakeholders want fast, real-time data about traffic lights, congestion, buses, and trains. They want to use the analytics to transform their business operations day-to-day. We didn't have that wealth of data.

— Roger Jones, Chief Technology Officer, Auckland Transport

Auckland Transport was formed in 2010, when the transport functions of the eight former Auckland local authorities and the Auckland Regional Transport Authority were combined, to oversee roads, traffic networks, and public transportation. The merger yielded five different operational centers with various technologies. A small staff monitored hundreds of older CCTV screens and tracked inputs on pedestrians, cyclists, and vehicles.

Making the roads safe necessitates pinpointing hot spots and trends, mitigating and reacting swiftly to issues, and monitoring the performance of the entire transportation network. In addition, multiple stakeholders and partners—from police and emergency responders to third-party application developers—need actionable insight on travel activities. The agency faced the challenges of launching a new CCTV system, converging the units and their data, and then assimilating and churning out vast data volumes to those who need to know. Auckland Transport selected a data analytics solution that enables personnel to derive insights and patterns from massive amounts of real-time streaming video data.

Putting the data to work, the agency has gained an integrated ticketing system with insight on travel times, patterns, trip frequencies, and demographics. Current statistics and other significant volumes of data, such as the parking system, reside on an analytics platform which processes structured data quickly. Auckland Transport can now remotely enforce traffic rules on special vehicle lanes with video analytics. This reduces operational costs and increases compliance with bus lane regulations. The agency's latest success is the well-publicized deployment of a special vehicle lane use enforcement system based on analytics, which, for the first time, provides a practical and automated mechanism to identify and generate evidence to fine vehicle owners for illegally using special vehicle lanes such as bus lanes. By providing a real deterrent to back up the ruling set by the New Zealand Transport Agency, Auckland Transport will promote safer and more efficient road use for all commuters in the city of Auckland.

Across sectors—from banking and manufacturing to healthcare, infrastructure, and government—the pattern is consistent: when agents are grounded in governed enterprise information and coordinated through orchestration layers with human oversight, workflows shift from reactive to anticipatory. The result is not incremental automation but a structural redesign of how organizations sense, decide, and act.

Agentic AI transforms fragmented systems into coherent operating models—embedding policy, accountability, and execution directly into the fabric of operations. Across every industry examined in this chapter, a consistent principle emerges: the lasting competitive advantage is not the agent performing the work—it is the governed evidence trail proving the work was done correctly, within policy, and under appropriate oversight. Enterprises that build governance into every workflow from day one gain a permanent structural advantage over those who attempt to retrofit it after deployment.

Executive Implications

CAIO

Regulated agent deployment requires governance aligned to jurisdictional standards. The CAIO ensures industry-specific AI regulations and enterprise governance frameworks are integrated so agents move from pilot to production safely and compliantly.

CIO

Agentic AI is an operating model shift, not a feature enhancement. Architecture must support orchestration across systems, enforce policy at the execution layer, and ensure agents are grounded in governed enterprise information—not isolated models bolted onto legacy workflows.

CFO

Value creation expands beyond cost reduction. Agentic workflows compress cycle times, improve working capital efficiency, reduce compliance exposure, and unlock margin through decision velocity—requiring updated ROI models that capture speed, resilience, and risk mitigation.

CHRO

As workflows become orchestrated, human roles evolve from task execution to supervision, exception handling, and policy design. Workforce strategy must focus on redefining accountability, strengthening AI literacy, and formalizing human-in-command escalation paths.

CDO

Agentic performance depends on defensible data foundations. Data lineage, quality controls, residency management, and governed access determine whether agents can reason safely and at scale across domains and jurisdictions.

COO

Operational excellence shifts from process optimization to workflow orchestration. Agents must reduce variance, anticipate disruption, and maintain continuity across value chains—turning resilience and responsiveness into structural capabilities rather than reactive interventions.

Now that you have seen agentic workflows in action across industries, the focus shifts from what agentic AI can do to how it is built, deployed, and governed. In the next part of this book, *Technology and Methods*, we examine the architectural foundations, deployment models, and agent lifecycle disciplines required to operationalize the agentic genome at enterprise scale.

Part 3

Technology and Methods

03

Chapter Eight

Reference Architecture for Agentic Enterprises

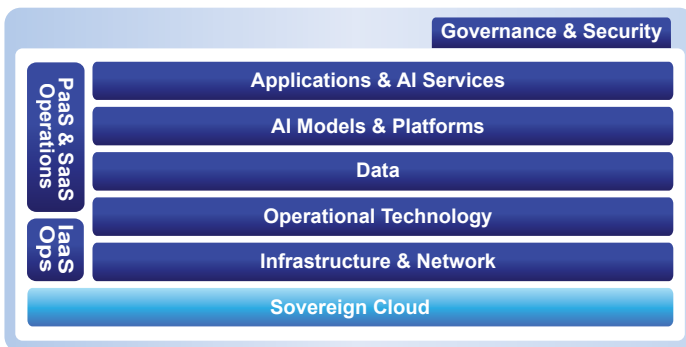
This chapter presents a practical reference architecture for agentic enterprises—one that binds content platforms as enterprise memory, agents as digital labor, orchestration as the system of action, and governance and trust as a continuous control plane, all grounded in sovereign deployment patterns that make scale defensible.

According to Forrester:

“ *Agentic AI systems are poised to not only become the backbone of the knowledge economy but will completely redefine how organizations operate and compete.*”¹⁴

The shift from GenAI to agentic AI marks a change in how enterprises design their operating model, along with how they deploy technology. Copilots and retrieval-augmented generation (RAG) improved access to information and accelerated individual productivity, but they didn't fundamentally change how work flows across systems. Agentic AI does. The moment AI systems can plan, coordinate, and act across multiple enterprise platforms, architecture becomes the determinant of value and risk.

This architectural framing is a direct evolution of the sovereign enterprise AI patterns introduced in the first book in this series, *Enterprise Artificial Intelligence: Building Trusted AI in the Sovereign Cloud*. In that book, we established trusted data, secure execution, and jurisdictional control as prerequisites for enterprise-grade AI. The difference now is that the cost of architectural weakness is higher. When systems can act, fragmented data, brittle integrations, and governance bolted on after the fact no longer just slow outcomes; they create operational and regulatory risk. The GenAI divide described in recent industry research—where substantial enterprise investment has produced uneven or negligible returns—has less to do with model quality than with the absence of platforms that connect intelligence to governed execution at scale.¹⁵

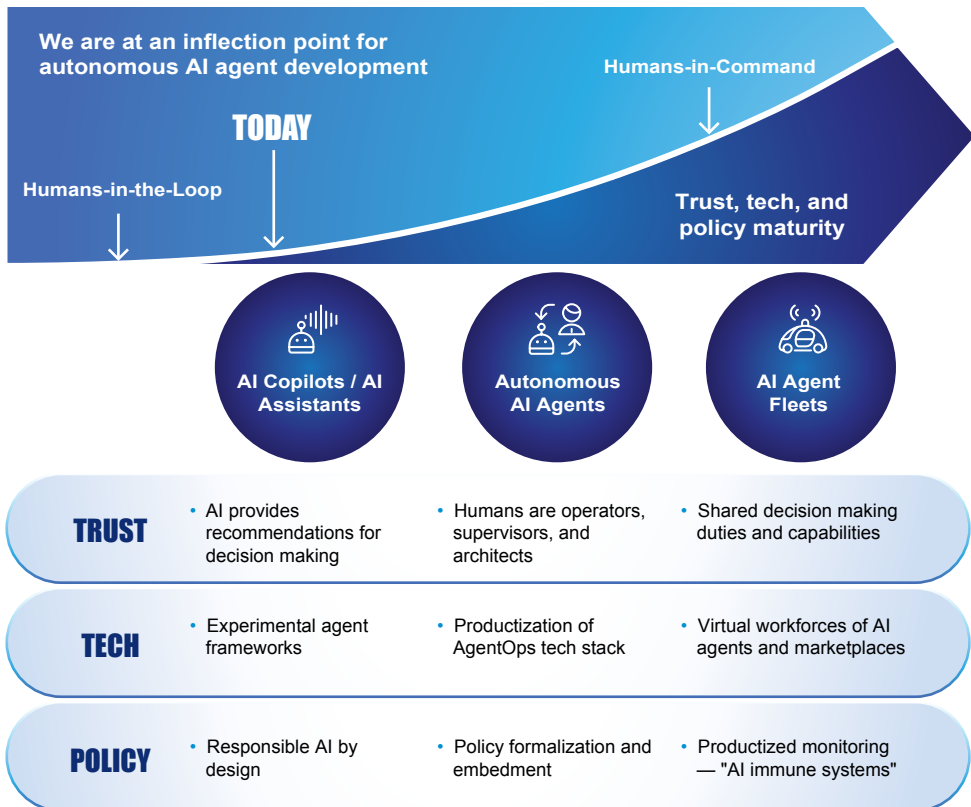


Enterprise Agentic Architecture

AI agents will evolve rapidly, progressing from task and application specific agents to agentic ecosystems. This shift will transform enterprise applications from tools supporting individual productivity into platforms enabling seamless autonomous collaboration and dynamic workflow orchestration.¹⁶

– Anushree Verma, Sr Director Analyst at Gartner®

The agentic enterprise is an operating model in which people define goals and trade-offs while digital agents execute coordinated work across functions.¹⁷ Research analysts across the board have projected rapid growth in enterprise applications that embed AI agents. They have even predicted the next phase: an emergence of “agent fleets,” where organizations orchestrate multiple specialized agents across end-to-end processes rather than deploying isolated AI features.

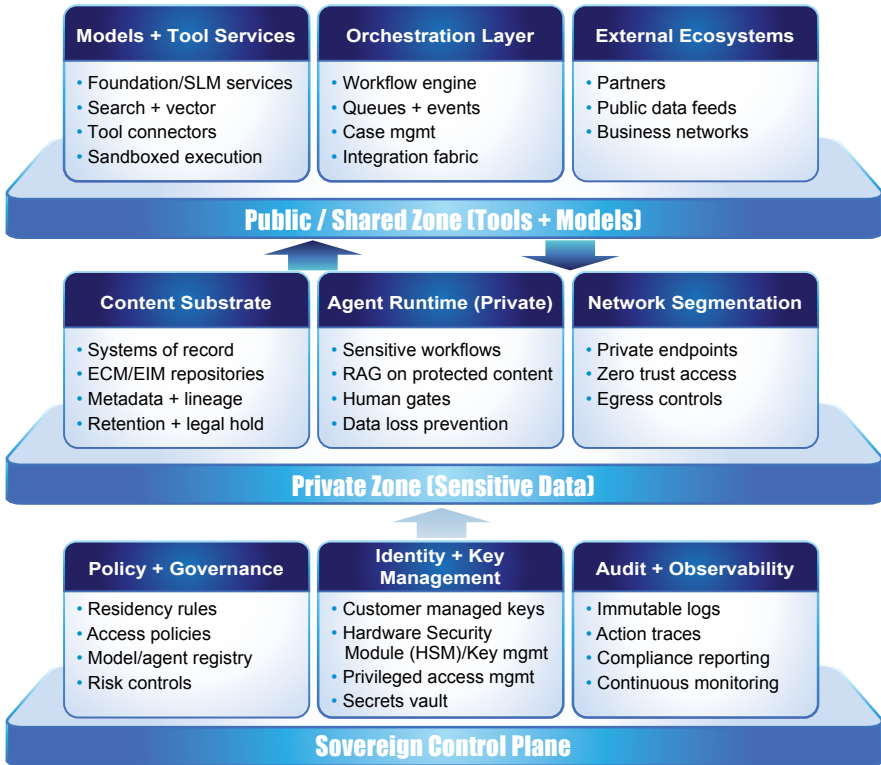


From AI Assistants to a Fleet of AI Agents¹⁸

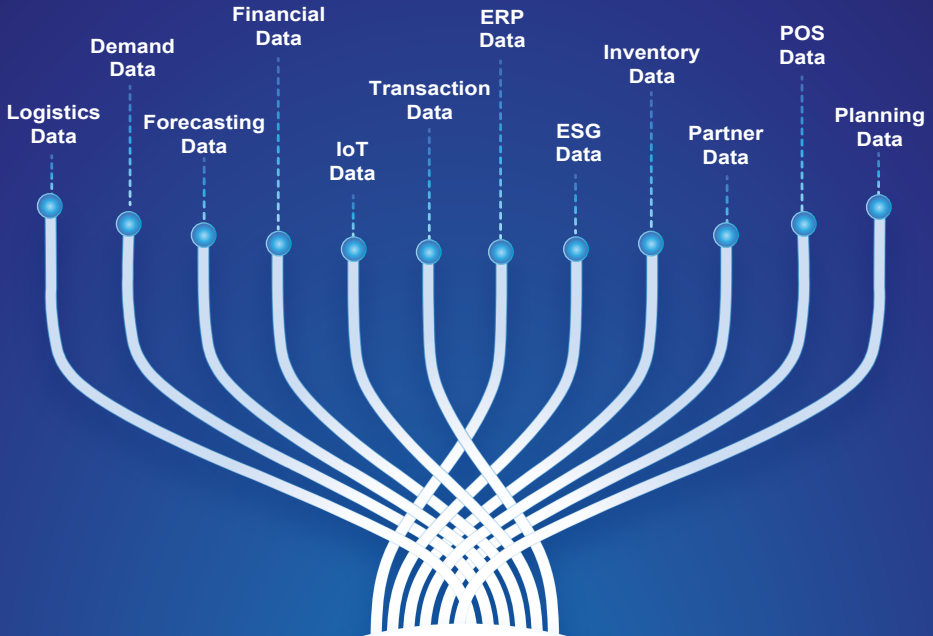
The common thread is that agentic AI becomes valuable when it is platformed: when orchestration, memory, and governance are designed as shared capabilities rather than stitched together per use case.

This is why reference architecture matters. A reference architecture is not a vendor blueprint; it is a repeatable pattern that helps CIOs, CDOs, CISOs, and enterprise architects make consistent decisions about where intelligence runs, what it is allowed to touch, how actions are governed, and how the organization proves it acted responsibly. In the agentic era, those decisions must be made together. The four layers presented in this chapter—content substrate, agent runtime, orchestration, and governance and trust—are not optional components. They are the minimum structure required to move from experimentation to durable, auditable execution.

The agentic enterprise architecture diagram shows how these layers interact, with sovereign cloud as the foundation that enforces jurisdictional control, key custody, segmentation, and auditability. The sovereign deployment model diagram (below) illustrates how sensitive enterprise memory and decision logs remain protected within jurisdictional boundaries while allowing controlled interaction with models and external tools through policy-enforced gateways.



Sovereign Model for Agentic AI Deployment



Data Feeds Every Business Process In A Company

Content Substrate – Content Platforms as Enterprise Memory

In agentic enterprises, content platforms stop being passive repositories and become the organization's memory. This memory is not limited to documents. It includes systems of record, business network transactions, operational telemetry, and the metadata that makes information governable: ownership, lineage, sensitivity, retention, and consent. Agents can only act responsibly when they can ground their reasoning in authoritative sources and understand the context and constraints attached to those sources.

The sovereign enterprise AI architecture in the first book emphasized zoning and control of enterprise information to preserve trust under regulatory scrutiny. That principle becomes even more important with agentic execution. When an agent updates a case file, triggers a payment workflow, or initiates a compliance action, the enterprise must be able to demonstrate what information was used, whether it was current and permitted for that purpose, and how the action aligned with policy. Data feeds every business process in the agentic enterprise. Without disciplined enterprise content management practices and metadata-rich content services, agents quickly inherit the same brittleness that plagued earlier automation initiatives—only at machine speed.

Practically, this means designing content platforms as services, not silos. Retrieval must support keyword, semantic, and vector search under policy, while event streams and immutable logs provide evidence for every instance of reading and writing. Records management, legal hold, and retention policies are not peripheral concerns; they define the boundaries within which agents can reason and act. In regulated sectors, the content substrate is the difference between scalable automation and non-defensible automation. All of this can be consolidated into a “single source of truth” with Enterprise Information Management, or EIM.

When enterprise memory shifts from a passive archive to an active data stream, it becomes the fuel for digital labor. This Lab Logistics Company illustrates that shift, transforming a manual ERP into an integrated content substrate where autonomous agents can orchestrate thousands of global orders with machine-speed precision.



A Leading Lab Logistics Company



A Lab Logistics Company supports a cooperative network of laboratory dealers across Europe, Asia, and Australia, managing inventory of more than 14,000 products and fulfilling thousands of monthly orders. Historically, the company relied on manual, paper-intensive processes tied to its legacy ERP system, requiring staff to enter and reconcile thousands of invoices and transactions by hand. Partnering with modern integration technology dramatically reduced paperwork, improved decision-making through system integration, and streamlined compliance activities such as ISO auditing, creating a more scalable foundation for growth.

In an agentic transport and logistics context, this foundation becomes the basis for autonomous orchestration of high-volume workflows. Instead of relying on manual invoice processing and human data entry, sensing agents could continuously monitor incoming documents, using intelligent classification and extraction to convert unstructured invoices into structured records. Compliance agents would automatically validate entries against business rules and supplier agreements, flagging only exceptions for human review. Orchestration layers then drive cross-system synchronization, updating finance, inventory, and partner systems without human handoff, while audit agents log every decision for traceability and governance. This approach accelerates end-to-end processing, reduces errors, and expands capacity beyond human limits.

By evolving from point automation to an agentic operating model, the company can support more dynamic partner relationships, improve responsiveness to demand shifts, and reallocate skilled staff away from repetitive work into strategic oversight. The agentic AI genome turns trusted content and system integration into a proactive engine of operational resilience, enabling real-time invoice throughput, compliance defensibility, and service delivery that keeps pace with market complexity.

Gartner predicts:

“ Up to 40% of enterprise applications will include task-specific AI agents by 2026, up from less than 5% today.¹⁹

”

Agent Runtime: Digital Labor with Tools, Memory, and Constraints

Agents are not chatbots. They are role-bound executors with goals, tool access, bounded memory, and constraints. An enterprise-grade agent runtime provides planning and reasoning loops, connectors to enterprise tools and workflows, governed short- and long-term memory, and comprehensive observability so that every action can be traced. The safety and security framework for agentic systems underscores why this matters: multi-step action chains, tool misuse, and emergent behaviors introduce risks that do not exist when AI is limited to text generation.

This is where architectural discipline turns into operational assurance. Least-privilege access must apply not only to people but to agents. Tool use must be constrained through allowlists, rate limits, and policy checks. High-impact actions require explicit human-in-command gates. Agentic AI eclipses humans-in-the-loop with humans-in-command, evolving from humans taking part in the decision-making process to humans defining the strategy, goals, and ethical boundaries of an agentic workflow—and the agent acting autonomously within those boundaries. Outputs destined for downstream systems should be structured to reduce ambiguity and brittleness. These controls do not slow value creation; they make it sustainable. Without them, early productivity gains are offset by incidents, rework, and regulatory exposure.

Orchestration: The System of Action

If content is memory and agents are digital labor, orchestration is motion. It is the layer that coordinates multi-agent workflows, manages dependencies and handoffs, integrates across systems of record, and enforces business policy in real time. Orchestration is what converts intelligence into outcomes.

This distinction explains why many GenAI pilots plateau. When intelligence is embedded only at the interface—answering questions or drafting content—value is constrained to individual tasks. When orchestration binds intelligence to end-to-end workflows—order-to-cash, procure-to-pay, incident-to-resolution—value compounds. Orchestration also provides the hooks for measurement: latency, cost, throughput, SLA adherence, and exception rates become observable properties of intelligent workflows, enabling CIOs to move from “we deployed a model” to “we reduced cycle time.”

Although early examples of AI agents are promising and offer autonomous decision-making, this technology still requires stronger accuracy, trust, and coordination to become mainstream.²⁰

In the agentic enterprise, people steer outcomes while agents execute, and this implies orchestration as the enterprise control plane. Agent fleets presuppose a coordination layer that disperses work across specialized agents and reconciles their output. The differentiation between assistants and task-specific agents stresses that execution requires workflow integration, not just conversational interfaces. Orchestration is therefore not an implementation detail; it is the enterprise operating system for agentic AI.

Governance and Trust: From Model Safety to System Safety

Traditional AI governance focused on data access and model risk. Agentic AI requires governance of behavior. The governance and trust layer provides the policies, controls, and evidence that make autonomous action auditable and reversible. Identity and access management must encompass agents as first-class actors. Policy engines must encode residency, consent, approvals, and segregation of duties. Model and agent registries establish what is approved for which tasks and under what constraints. Continuous evaluation and monitoring surface drift, safety issues, and compliance deviations before they become incidents.

Organizations must establish clear data governance policies that ensure agentic AI systems have access to high-quality, current information while maintaining appropriate security controls. Governance frameworks should define clear boundaries for autonomous operation, escalation procedures for edge cases, and accountability structures that maintain human oversight over strategic decisions.

The research on agentic system safety emphasizes that trust is an emergent property of systems, not models. Governance-by-design is therefore essential. It moves trust from being an aspiration to being an architectural property. It must be operationalized, not managed as a side committee. When governance is embedded in orchestration and agent runtime, enterprises can scale intelligent execution without sacrificing control.

Trust in an agentic system is an architectural property, not an aspiration. The following Global HR Company case study demonstrates how a fragmented document environment was transformed into a policy-aware execution layer, where specialized agents now autonomously enforce jurisdictional labor laws and privacy controls by design.



A Global HR Company



A global enterprise operates a complex HR environment supporting employees across multiple regions, each governed by distinct labor laws, privacy requirements, and internal policies. Historically, HR document management—contracts, onboarding forms, benefits records, and compliance documentation—relied on fragmented systems and manual processes. This created delays in onboarding, inconsistent application of policies, and elevated audit and compliance risk.

By modernizing HR document management on a centralized content platform, the organization standardized records, automated routing and approvals, and enforced retention and privacy controls. This delivered immediate efficiency gains by reducing cycle times, improving visibility, and strengthening compliance. However, these workflows still depended heavily on human intervention to interpret policy changes, validate compliance across jurisdictions, and resolve exceptions.

In the agentic era, this content foundation becomes a policy-aware execution layer for HR operations. An agentic HR orchestrator continuously monitors employee lifecycle events and relevant content, while specialized agents validate documents against jurisdictional labor laws, internal policy rules, and privacy requirements before issuance. Classification and retention agents automatically apply the correct access controls and records policies, while exception agents escalate anomalies to HR stewards with full context and audit trails. Routine tasks—document creation, verification, routing, and archival—are executed autonomously, with governance enforced by design.

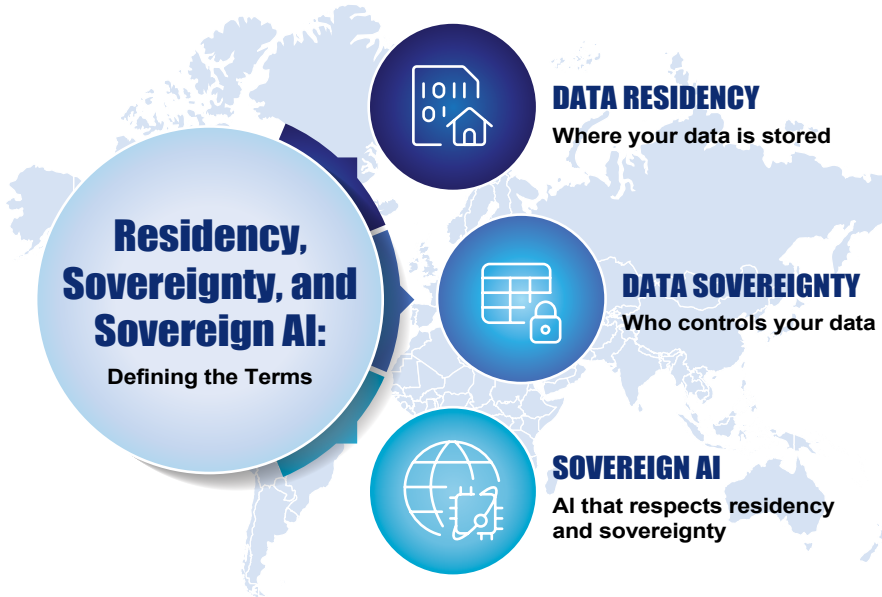
The result is materially higher efficiency without compromising trust. Onboarding cycles shorten, compliance errors decline, and audit readiness improves because every automated action is traceable and policy aligned. HR teams are freed from administrative burden to focus on strategic workforce priorities, while leadership gains confidence that automation operates within regulatory and governance boundaries.

Sovereign Cloud: The Control Boundary for Scale

Sovereign cloud is often misunderstood as a geopolitical posture. In practice, it is the mechanism by which enterprises prove control over data, models, keys, and execution under regulatory scrutiny. The sovereign EAI architecture presented in the first book in this series established segmentation, customer-managed keys, data residency, and auditable gateways as foundational. Those principles become even more consequential in agentic systems, where agents may write back to systems of record, trigger financial or legal workflows, and coordinate across partners.

Recent research highlights the pressure enterprises face as AI adoption accelerates while regulatory expectations intensify. The *State of AI in Business* report from MIT underscores the uneven returns organizations are seeing despite significant investment, reinforcing the reality that scale without architecture and governance erodes value.²¹

Sovereign patterns—private zones for sensitive content, governed interfaces to shared services, and immutable audit trails—provide the structural assurance needed to move from pilots to platforms. They enable hybrid and multi-cloud deployment without losing jurisdictional control, ensuring that intelligent execution remains defensible across borders and business units.



Clarifying the Language: Residency, Sovereignty, and Sovereign AI

The sovereign deployment model diagram in this chapter illustrates how policy and key management in a sovereign control plane govern interactions between private zones containing enterprise memory and shared zones hosting models and tools. This pattern preserves agility while maintaining compliance and evidentiary rigor.

Putting the Architecture Together

The sovereign cloud foundation defines where data lives, how keys are controlled, and how networks are segmented. The content substrate provides authoritative enterprise memory with the metadata and lineage required for governed reasoning. The agent runtime supplies digital labor with constrained tools and traceable actions. The orchestration layer coordinates multi-step execution across systems and agents. The governance and trust layer enforces policy, monitors risk and preserves auditability.

The reason these layers must be designed together is simple. If any one is missing, the system either cannot scale or cannot be trusted when it does. Architecture is the difference between intelligence that demos well and intelligence that operates under scrutiny.

Enterprise agentic AI success depends on robust technical infrastructure that can support autonomous operation while maintaining security, compliance, and performance standards. This includes scalable computing resources, comprehensive monitoring and logging systems, and integration frameworks that enable seamless interaction with existing enterprise applications.

Executive Implications

As agentic AI matures, standardized protocols and frameworks will enable seamless interoperability, allowing agents to sense their environments, orchestrate projects and support a wide range of business scenarios.²²

— Anushree Verma, Sr Director Analyst at Gartner®

The competitive frontier is no longer infrastructure alone; it is orchestration at machine speed. Speed without control, however, is volatility. The agentic enterprise wins by combining governed enterprise memory, coordinated execution, scalable digital labor, and provable control grounded in sovereign deployment. This is how agentic AI becomes an enterprise operating system—one that moves faster without being reckless and automates without being opaque.

As agentic AI becomes embedded in core operations, each executive role inherits new architectural responsibilities that shape value, risk, and control.

CAIO

Regulated agent deployment requires governance aligned to jurisdictional standards. The CAIO ensures industry-specific AI regulations and enterprise governance frameworks are integrated so agents move from pilot to production safely and compliantly.

CIO

Moving from pilots to platforms requires an orchestration-first architecture that integrates agents, content, and systems of record into governed, end-to-end workflows.

CFO

Agentic AI only produces defensible ROI when autonomous execution is grounded in auditable workflows, sovereign deployment patterns, and cost-visible orchestration rather than isolated AI features.

CHRO

As digital labor scales, HR must govern human-agent collaboration through clear escalation paths, role boundaries, and accountability frameworks embedded into the operating model.

CDO

Enterprise memory becomes operational infrastructure; without metadata-rich content platforms and lineage-aware data services, agentic execution amplifies risks in data quality, compliance, and trust.

COO

Sustainable cycle-time reduction depends on treating orchestration as the system of action, enabling coordinated multi-agent execution across functions rather than speeding up individual tasks.

The architecture defined in this chapter becomes actionable only when it is deployed across hybrid and sovereign environments and sustained through disciplined lifecycle management of agents—topics we turn to next. In the following chapter, we'll examine the deployment models for agentic AI.

Chapter Nine

Deployment Models: Multi-Cloud, Hybrid, & Sovereign

Agentic AI changes the deployment conversation because it changes the workload. A chatbot can live on the edge of the enterprise and still be useful. An agentic system cannot. Agents don't just generate text; they read governed information, call tools, trigger workflows, write back to systems of record, and leave audit trails behind them. That means your deployment model is no longer a pure IT preference. It becomes a business risk decision, a compliance decision, and increasingly, a national competitiveness decision. Orchestration at machine speed only works when deployment, data residency, identity, logging, and governance are designed as one system.

In this chapter, we examine the agentic enterprise deployment models, including multi-cloud, hybrid, and sovereign across jurisdictions.

Why “Deployment” Is a Board-Level Topic in the Agentic Era

In classical enterprise AI, deployment was largely about where models run and where data is stored. In agentic AI, deployment is about where decisions are executed. Once agents can initiate actions across finance, HR, supply chain, safety, legal, or citizen services, the enterprise has to answer uncomfortable questions:

Where does the agent run, and where do its tools live?

Where does it retrieve memory and evidence from?

Where is the audit log created, stored, and reviewed?

Who can inspect the chain-of-thought or the decision trace, and under what legal authority?

What happens when an agent crosses borders—even unintentionally—by calling a service hosted in another jurisdiction?

Agentic AI encompasses institutional memory plus reasoning at enterprise scale—and in sovereign contexts, at national scale. Your deployment model either preserves that memory and defensibility, or erodes it.

Sovereignty and control are becoming a first-class architectural requirement. Research analysts have explicitly framed the market shift toward sovereign cloud platforms as a distinct evolution in how digital sovereignty is implemented. The sovereign cloud is not just a “region setting.” It is shaped by jurisdictional control, compliance, and operational constraints.

Agentic AI forces this shift into daylight because, once agents can act, you must be able to prove what happened and why. And proof is a deployment outcome: it depends on where logs live, where keys are controlled, where retention is enforced, and where policy decisions are executed.

Crossing the GenAI Divide

Industry research on the GenAI Divide adds a crucial nuance: many AI efforts stall not because of budget or excitement, but because tools don’t learn, don’t integrate, and don’t fit workflows. The report on the topic from MIT describes a steep pilot-to-production drop-off for task-specific enterprise GenAI tools and calls out that only a small share reach production at scale.²³

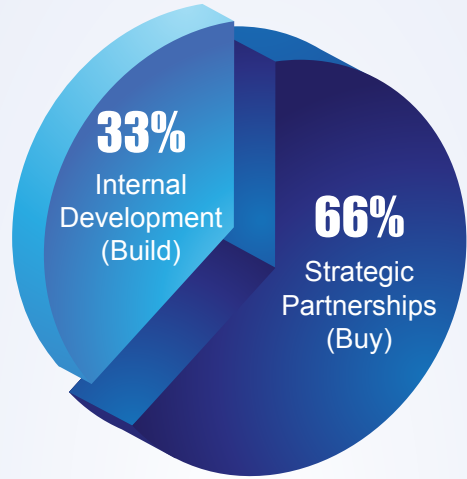
It also makes a procurement point that matters directly to deployment strategy: organizations that cross the divide act like Business Process Outsourcing (BPO) clients, not SaaS customers. They demand deep customization, push adoption from the front lines, and hold vendors accountable to business metrics—because crossing the divide requires partnership rather than purchase.

This is exactly where deployment models become strategic. If you treat agentic AI like a generic SaaS add-on, you will select deployment for convenience. If you treat it like an operating capability with workflow accountability, you will select deployment for control, evidence, and integration depth.

As illustrated in the graph on this page, vendor partnerships materially outperform purely internal builds in reaching deployment. For CIOs and executives, the practical implication is not “always buy.” It is that build-buy decisions should be driven by who can help you operationalize governed workflows fastest—because agentic ROI appears when systems are integrated, not when demos are impressive.

Crossing the GenAI Divide requires treating AI as a core operating capability rather than a generic SaaS add-on.

By automating complete application stacks, the Logistics Group in the following case study has built the deep integration and metadata-rich environment necessary to move from simple task automation to autonomous, agent-driven operations.



% of Deployments

Strategic Partnerships (Buy) — 66%

Procure external tools, co-develop with vendors

Internal Development (Build) — 33%

Build and maintain GenAI tools fully in-house

Hybrid (Build-Buy) — Insufficient Data to Quantify

Internal team co-develops with an external vendor

Research Limitations: These percentages reflect an interview sample of 52 organizations and may not represent broader market patterns. Success definitions varied across organizations, and observation periods may not capture long-term implementation trends.

Three Structures for GenAI Deployment²⁴



A Logistics Group



Faced with long lead times and manual provisioning for its core ERP management system, a diversified logistics group adopted cloud service automation and operations orchestration to help them eliminate repetitive manual tasks and accelerate environment provisioning across its hybrid infrastructure. The end result was more consistent, scalable deployments.

Rather than relying on handcrafted scripts and human-driven checklists, the company automated the end-to-end setup and configuration of complete application stacks. This not only shortened deployment cycles and reduced risk of configuration drift but it also created a rich operational telemetry fabric—a prerequisite for higher-order AI capabilities. With cloud-native orchestration and consistent metadata flowing from automated processes, the organization now has a trusted, governed operational context in which advanced agents can begin to act.

In an enterprise where both speed and compliance matter, this modernization positions the company to evolve beyond basic task automation into agent-driven operations—where AI agents can autonomously detect opportunities (e.g., bottlenecks or anomalies), recommend optimizations, and trigger orchestrated remediation steps, all anchored to governed cloud services and secure information flows.

A Simple Framework: Place Four Things, Not One

In Chapter 8, we framed agentic architecture as five layers. When we discuss deployment models, it helps to collapse the conversation into four placement decisions that recur across every enterprise:

1. **Compute:** where the models and orchestration execute
2. **Memory:** where the content substrate lives (and where retrieval is performed)
3. **Tools:** where connectors and action endpoints live (ERP writes, case updates, workflow triggers)
4. **Trust:** where identity, policy enforcement, logging, and audit live

Multi-cloud, hybrid, and sovereign models are simply different solutions that address these four decisions.



Deployment Models for Agentic AI (Multi-Cloud, Hybrid, Sovereign)

The figure above illustrates how the five-layer Agentic Enterprise Architecture is deployed across public cloud, private cloud, sovereign cloud, and edge/on-prem environments (as part of a hybrid deployment). While elastic compute and non-sensitive inference can run in shared cloud environments, governed content, privileged agents, and audit trails are anchored in private and sovereign zones. The model demonstrates how enterprises balance agility with jurisdictional control, ensuring that agentic systems can scale without compromising trust, compliance, or evidentiary integrity.

Model 1: Multi-Cloud – One Architecture, Multiple Clouds

Multi-cloud is often misunderstood as “we use more than one cloud.” In successful enterprises, multi-cloud is a risk and resilience posture: multiple clouds are used to manage vendor dependency, service fit, latency, regional availability, M&A inheritance, and regulatory constraints. In the agentic era, multi-cloud also becomes a workload placement tool: different parts of the system can run in different environments, as long as trust boundaries remain intact.

A practical pattern is to treat public hyperscalers as elasticity providers—excellent for burst compute, simulation environments, non-sensitive experimentation, and scalable inference—while keeping governed memory and systems-of-record integration under stricter control. While hyperscaler may be used for elastic training compute, non-sensitive analytics, and simulation, they should not be relied on as systems of record. The “sovereign spine” must remain sovereignly controlled.

That “spine” concept generalizes cleanly to multinationals: the enterprise spine is the governed content substrate and trust layer that anchors identity, policy, auditability, and records. Once this is made explicit, multi-cloud stops being messy and becomes modular.


Your cloud strategy must be aligned with risk, compliance, and operating continuity—especially as AI increases cross-system automation. McKinsey has emphasized “futureproofing” IT functions amid disruptions in ways that typically prioritize resilience, security, and operating model readiness rather than single-vendor purity.²⁵

The key architectural guardrail is this: even if some model endpoints run in a public cloud, governed access to private content, tool permissions, and audit logs must remain enforceable. Otherwise, you have built an agentic system that can act but can't be proven safe.

Model 2: Hybrid – Where the Enterprise Actually Lives

Hybrid deployment is the default state for most large organizations because reality is hybrid. You have legacy systems, regulated data, acquisitions, on-prem operational technology, and cloud-native digital experiences—all at once.

In agentic AI, hybrid is the mechanism by which agents can see across the enterprise without violating boundaries. Hybrid architectures also make it possible to adopt agentic patterns incrementally: you can deploy the content substrate and governance layer first, integrate a few orchestrations, and scale agent by agent without forcing a “big bang” migration of every system.



The cloud is especially vital here because it provides the control plane patterns for orchestration, eventing, observability, and secure connectivity. But hybrid also forces discipline: you must decide which zones can be written to by agents, which zones can only be read, and where human approval gates are mandatory.

We framed this in the first book in this trilogy: AI can take accountable action when it is connected to governed information and can respect permissions, retention schedules, and regulatory requirements. Hybrid deployment makes that possible at scale if (and only if) policy enforcement is consistent across zones.

Model 3: Sovereign Deployments – In-jurisdiction Control as a Design Requirement

Sovereign cloud is not a region setting. It is an architectural commitment to jurisdictional control, key custody, auditability, and defensibility.

Sovereign deployments are often described narrowly as “data residency.” In practice, sovereignty is broader: it includes jurisdictional control of data, keys, access, operations, and legal authority—and it typically includes defensibility requirements (auditability, retention, evidentiary integrity).

If a national security infrastructure is required, a country cannot outsource jurisdictional control. This is the public-sector version of a pattern that regulated industries also recognize: in finance, healthcare, energy, and critical infrastructure, the question is not only “*where is the data?*” but “*who can compel access, who can operate the platform, and how is evidence preserved?*”

This is also why “sovereign evidence rails” become central. Paperless trade depends on immutability, replayable compliance, trusted evidence across agencies, and enforceable governance. This is exactly the governance-and-trust layer discussed in Chapter 8, expressed as an operational mandate. Sovereign cloud is emerging as a distinct category of cloud platform evolution, shaped by sovereignty requirements rather than generic cloud economics.

Cloud is not just where agentic AI runs. It is the control plane that enables elastic orchestration, real-time governance, and jurisdiction-aware placement of intelligence.

Why Cloud Is Vital for Agentic AI (and Why It's Not Sufficient)

Cloud is vital for agentic AI for five practical reasons:

1. **Compute elasticity.** Agentic systems are bursty: orchestration spikes during events, incidents, month-end close, border surges, or disruptions. Cloud provides elastic capacity without provisioning delays.
2. **Resilient orchestration primitives.** Event-driven architectures, managed queues, workflow runtimes, and observability stacks are the “plumbing” that make multi-agent coordination reliable.
3. **Consistent security building blocks.** Identity services, key management, secrets vaulting, policy engines, and audit pipelines are mature in cloud environments—and agentic AI is security-hungry because agents touch many tools.
4. **Geographic placement.** Multi-region and in-jurisdiction hosting is the only practical way to meet latency and sovereignty constraints simultaneously for global organizations.
5. **Operational tempo.** Agentic AI moves fast; cloud operating models (when governed) allow you to deploy controls, update policies, rotate keys, patch dependencies, and monitor drifts at a cadence that matches machine-speed workflows.

Cloud gives agentic systems elastic compute, managed security primitives, resilient orchestration patterns, and the ability to place capabilities close to where data is legally allowed to live. But cloud alone does not solve the hard part: trust. Trust comes from governed access to private, permissioned information and from enforcement of policy and auditability across every action an agent takes. Value comes when AI operates inside the firewall with governed access to private content, permissions, and policy-rich metadata, so it can act accountably and leave a verifiable trail.

Sovereignty and trust are governance outcomes. A cloud region does not automatically give you evidentiary integrity. A managed service does not automatically give you audit defensibility. You still have to design the system so every action is policy-governed, permissioned, logged, and reviewable.

Without private, permissioned data and the governance to use it responsibly, AI hits a ceiling. It can summarize the internet, but it can't safely approve an invoice or resolve an exception inside your systems. Agentic AI raises that ceiling, but only if deployment makes trust enforceable.

Agentic AI changes the deployment question from “Where do we run models?” to “Where do decisions execute, and where is evidence preserved?”

Agentic Deployment: New Safety and Security Requirements

Deployment models also reflect a core reality: agentic systems create different risk surfaces than traditional GenAI. The safety and security literature warns about risks such as tool misuse, cascading action chains, and emergent behaviors—which become more relevant as agents gain access to enterprise tools and can take multi-step actions.

In practical terms, this means deployments must support:

- Strong identity and least-privilege access at the tool layer
- Policy enforcement at runtime (not only at design time)
- Immutable, queryable audit logs for every read/write/action
- Human approval gates for high-stakes decisions
- Segmentation between public inference surfaces and private action surfaces
- Kill switches, rate limiting, and anomaly detection for agent behavior

These are not optional features. They are the minimum viable controls for scaling from pilots to an operating model.

A resilient deployment model must do more than just house data; it must provide a platform for continuous validation and auditable evidence. In the following case study, Velliv, Denmark’s leading pension provider, illustrates this by leveraging a hybrid cloud architecture to transform its quality assurance from a manual quarterly bottleneck into an intelligent, centralized automation framework that ensures every system update is governed and verified in real-time.

Velliv – Denmark’s Leading Pension and Life Insurance Company



The beauty of an automated test framework... is the full dashboard demonstrating to IT management testing efficiency and coverage.

– Benjamin Bennike Aagren, Test automation manager, Velliv

Velliv, a leading Danish pension provider with over 420,000 customers, must meet rigorous quality and regulatory compliance standards. Critical systems—including core pension, customer handling, ERP platforms, and customer-facing applications—require extensive testing to ensure reliability and pass regular audits. Test automation was almost non-existent, and regression testing of major releases was a huge manual workload, only done once every three months—limiting agility and increasing manual workload. To overcome these challenges and accelerate delivery without sacrificing quality, Velliv turned to an AI-powered functional testing tool to bring intelligent automation into the fold.

Velliv has implemented a centralized automation framework which has transformed its release cadence. Previously, regression tests were only run at the end of quarterly major release deployments; the company now runs their regression test at every deployment, multiple times a week in test environments. Major releases still occur every three months. However, with automated validation happening continuously, defects are caught much earlier. Much more is tested with much less manual effort. The quality assurance team can now focus on testing new features while their regression checks run in one click, ensuring ongoing reliability and readiness for every deployment to a test environment. The solution operates within Velliv’s private cloud, supporting security, flexibility, and scalability.

Velliv’s transformation proves that the path to agentic maturity begins with automated trust. By anchoring their AI-powered testing within a private cloud, they have ensured that their core pension systems remain compliant while drastically increasing the velocity of their deployment cycles. For regulated enterprises, this hybrid approach provides the necessary sovereign spine to scale intelligent operations without compromising the evidentiary integrity required for national and industry audits.

Putting it together: Deployment Patterns That Scale Securely

Across enterprises, regulated industries, and sovereign contexts, the most durable deployments converge on baseline requirements:

1. **A sovereign or controlled “spine” with elastic “limbs.”** Use sovereign/private zones for governed content, records, audit logs, and privileged tools; use public cloud for elasticity where appropriate; use edge/on-prem for local systems and operational environments. Hyscaler elasticity is useful, but not as the system of record.
2. **Separate retrieval from action.** Many organizations can tolerate broader retrieval (with strong controls) than they can tolerate autonomous writes. Deployment should reflect that: keep write-path tools and approvals closer to the governed core.
3. **Make audit and retention first-class infrastructure.** The EIM backbone is what makes actions defensible. If you cannot prove what happened, you will eventually be forced to slow down. This is the hidden failure mode of pilot success: pilots can work without defensibility; operating models cannot.
4. **Treat adoption as an operating model change, not a product rollout.** The concept of the GenAI Divide is explicit: the organizations that cross the divide drive adoption through distributed experimentation, vendor partnerships, and clear accountability. Deployment models should enable this by providing safe sandboxes, governed pathways to production, and measurable outcome instrumentation.

Deployment: Key Challenges

Forrester forecasts that by 2030, spending on off-the-shelf AI governance software will more than quadruple, reaching \$15.8 billion and capturing 7% of overall AI software spending. This trend reflects the urgency for organizations to manage AI's integrity amid rapid AI adoption and increasing regulatory demands.²⁶

At the early stages of the agentic AI revolution, one lesson is clear: it takes real work to do this well. While a small number of organizations are beginning to see productivity gains from agentic deployments, many more are struggling to move beyond pilots into durable operational value. This pattern is not unusual. Every major technology shift follows a similar arc of early enthusiasm, uneven execution, and eventual maturation. The difference with agentic AI is that the consequences of failure are operational. When agents are embedded in real workflows, shortcomings surface quickly.

One of the most consistent lessons emerging from early deployments is that success is rarely about the agent itself—it's about the workflow. Organizations that focus primarily on building impressive agents often find that overall process performance barely improves. The systems may look capable in isolation but fail to deliver meaningful outcomes across end-to-end business processes. The organizations that generate real value begin by reimagining workflows across people, processes, and systems. We explore this in greater detail in Chapter 11, From Pilots to Platforms.

In mature deployments, agentic systems are designed to learn within the workflow itself. User edits, corrections, overrides, and approvals become feedback signals that refine agent behavior over time. This creates a self-reinforcing system: the more agents are used in real operational contexts, the more aligned and effective they become. Without these learning loops, organizations risk deploying agents that stagnate while the business evolves around them.

At the same time, leaders are learning that agents are not always the right tool for every job. Not all workflows benefit from autonomous behavior. Highly standardized, tightly governed processes often perform better with rules-based automation, deterministic decision engines, or traditional analytics. In contrast, high-variance workflows that require interpretation, synthesis, and judgment—such as complex financial analysis, investigative casework, or multi-party coordination—are better candidates for agentic approaches. The most effective deployments avoid a binary “agent or no agent” mindset and instead assemble the right mix of tools for each step of the workflow under a common orchestration framework.

Gartner states,

AI transformation is being built on AI governance. And predicts that, By 2027, fragmented AI regulation will grow to cover 50% of the world's economies, driving \$5 billion in compliance investment.²⁷

Trust remains one of the hardest deployment challenges of agentic AI. Many organizations report strong demo performance but poor user acceptance in production environments. Users quickly disengage when outputs feel low-quality, inconsistent, or difficult to verify. Over time, this erosion of trust can negate any efficiency gains from automation. High-performing teams treat agents less like software features and more like new employees. They define clear roles for agents, establish performance expectations, and invest in structured evaluation frameworks that test outputs against expert judgment. Continuous evaluation becomes a core part of deployment, not an afterthought.

Deployment in the agentic enterprise is not merely a technical exercise. It is an operating model transformation—one that reshapes how people, systems, and intelligent agents work together to deliver outcomes at machine speed, without sacrificing trust, accountability, or human judgment. These aspects of the agentic enterprise are discussed in Part 4 of this book on the execution and adoption of AI.

If you want a simple summary for executives, it's this: agentic AI is a deployment problem before it is a model problem. The difference between “good demos” and defensible automation is not prompt engineering. It is architecture—and deployment is where architecture becomes enforceable.

Executive Implications

CAIO

Multi-cloud and sovereign deployments increase governance complexity. The CAIO ensures AI deployments remain compliant across jurisdictions by aligning enterprise governance standards with regional regulatory requirements.

CIO

Hybrid, multi-cloud, and sovereign deployment models must be designed as a single control fabric, not separate environments, to preserve consistency, security, and operability at scale.

CFO

Uncontrolled cloud sprawl and duplicated AI stacks erode the economics of agentic AI; disciplined deployment patterns and shared platforms are essential to sustain margin and predictability.

CHRO

Distributed deployment increases operational complexity for teams; workforce enablement must include cloud literacy and clear ownership models for hybrid agentic operations.

CDO

Data residency, cross-border flows, and model access policies must be embedded into deployment architecture, or regulatory exposure will scale as fast as AI adoption.

COO

Operational resilience in agentic systems depends on deployment architectures that preserve continuity of service across clouds and jurisdictions without fragmenting workflows.

Agentic AI readiness is not about how many agents are deployed. It is about whether the enterprise can govern, observe, and integrate autonomous action into real workflows—safely, at scale, and across jurisdictions.

In the next chapter, we'll take a closer look at the agent lifecycle, including training, monitoring, evolution, and maturity at scale.

Chapter Ten

The Agent Lifecycle

The Agentic Development Lifecycle (ADLC)—distinct from the traditional Software Development Lifecycle (SDLC)—is emerging as a specialized lifecycle for building, deploying, and governing autonomous agents at enterprise scale. It is an iterative framework for creating systems that evolve over time—moving beyond static prompt-based models into continuous behavioral quality management, tool integration, and autonomous decision-making under uncertainty.

In this chapter, we'll examine all five stages of ADLC, as well as the fundamental challenges and best practices for building, deploying, and governing autonomous agents at scale.

What ADLC Optimizes For (That SDLC Does Not)

Agentic AI is more than “software with a prompt.” It is an operating system for goal-directed work: systems that interpret intent, plan, call tools, take actions, and learn from outcomes. That shift breaks many assumptions baked into the traditional Software Development Lifecycle (SDLC). The SDLC is optimized for deterministic logic—code that should behave the same way every time. Agentic systems are probabilistic: their behavior varies with context, model updates, tool availability, data quality, and the evolving state of the world they act within. Treating agent delivery like SDLC is “dangerously obsolete,” and the right orientation is “evaluation-first” rather than “code-first.”²⁸

Agentic systems introduce three realities that fundamentally reshape lifecycle thinking:

- 1. Behavior is the product.** You are not shipping “features” as much as you are shaping *how the agent behaves* across a spectrum of real-world situations. That means intent, boundaries, and failure modes must be defined early and continuously re-validated.
- 2. Tools and integration are first-class.** Agents only create enterprise value when they can safely act through APIs and systems of record—within strict authorization boundaries and auditable workflows.
- 3. Governance is not a phase; it underpins everything.** In agentic systems, governance must be embedded across ideation, build, test, deployment, and post-launch observation—because the system keeps operating and changing in production.

Traditional SDLC practices assume deterministic logic and repeatable outcomes, while agents operate under uncertainty and can drift over time. As a result, testing shifts from primarily unit and functional tests to behavioral drills and alignment checks; control changes from fixed code paths to guardrails and approvals that constrain autonomy; and maintenance evolves from periodic bug fixes to continuous refinement of behavior, tools, and oversight as goals and environments change.

An “agent control plane” is the governance and coordination layer of the agentic enterprise that defines, monitors, constrains, and audits how agents operate across workflows. This framing underscores the architectural implications of the differences between ADLC and SDLC. As agents proliferate across development, orchestration, and operational layers, governance cannot remain embedded only within individual applications. Instead, enterprises require an oversight plane that provides consistent visibility, policy enforcement, and risk management across a heterogeneous estate of agents, or the full managed enterprise “population” of agents. This architectural separation reflects the reality that agentic systems introduce new forms of autonomy and risk that demand centralized, lifecycle-spanning governance rather than localized, application-by-application controls.

To move from the architectural control plane to real-world velocity, we must look at how the principles of ADLC—behavioral consistency, tool integration, and continuous governance—manifest in high-stakes environments like digital banking, where the transition from deterministic code to probabilistic agentic behavior is already yielding measurable competitive advantages—as illustrated in the following case study.



A Leading Bank in Brazil



A leading Brazilian bank with over 19 million digital customers saw an opportunity to transform how it delivers software by adopting an agentic development lifecycle mindset focused on outcomes, automation, and continuous improvement. Facing rising customer expectations and competitive urgency, the bank moved to a unified, behavior-oriented delivery model in which tooling, quality practices, and operational telemetry are tightly integrated to support systems that learn and adapt over time. As part of this initiative, they implemented a functional testing lab for mobile and web to unify planning, quality, and execution across its agile pipelines.

This operating model aligns with the ADLC's core stages of ideation, composition, testing, deployment, and ongoing optimization. By anchoring delivery around measurable outcomes and continuous behavioral evaluation—rather than isolated task completion—teams were able to standardize workflows, reduce friction between development and quality functions, and accelerate release cycles while maintaining control and traceability. With hundreds of projects and workspaces operating daily, the bank's delivery ecosystem now functions as a continuous improvement loop, where performance signals from production feed directly back into design and testing decisions, reinforcing alignment between autonomous system behavior and business intent.

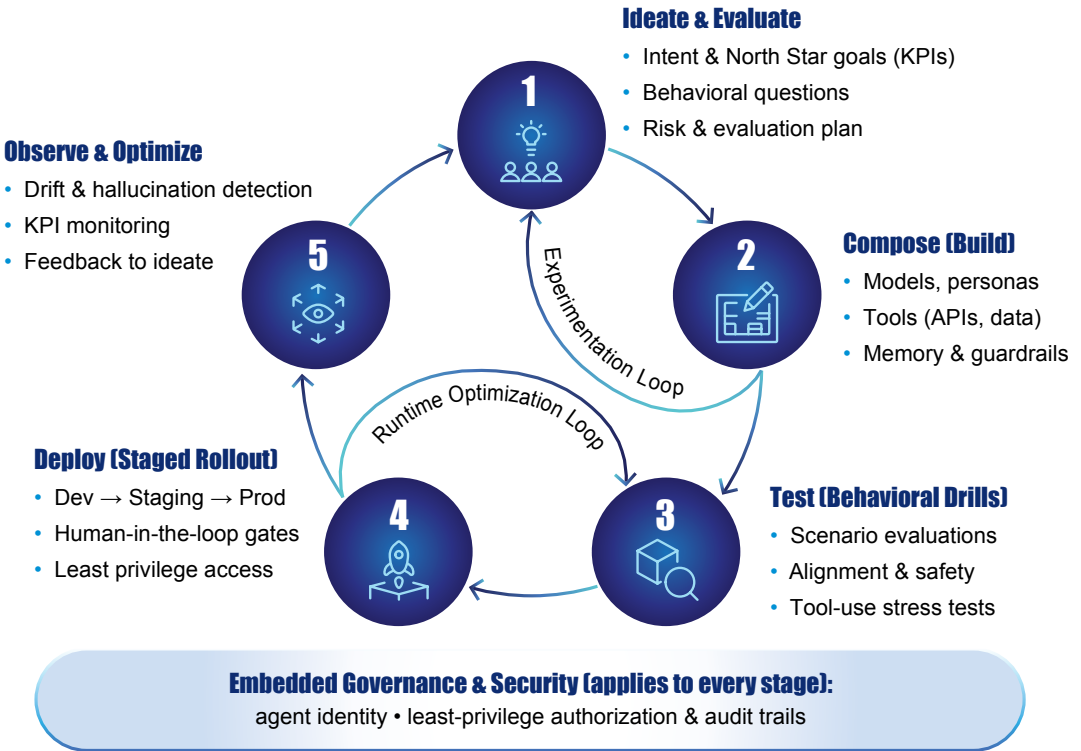
The Brazilian bank's transformation illustrates that the ADLC is not merely a technical upgrade, but a fundamental shift in the enterprise operating system. By treating delivery as a continuous behavioral loop rather than a linear track, the bank successfully navigated the core challenge of agentic systems: maintaining rigorous governance without sacrificing the autonomy required for speed. Their journey confirms that when behavior is the product, the lifecycle must prioritize the alignment of autonomous actions with business intent, ensuring that as the estate of agents grows, so too does the organization's capacity for reliable, scaled innovation.

The Five Core Stages of ADLC (Iterative Loop)

Multiple modern frameworks for ADLC converge on a five-stage, iterative lifecycle. Though they are often labeled with slight variations, they are consistent in intent: ideate and evaluate, build, test, deploy, and observe and optimize.

What follows is an enterprise-ready ADLC lifecycle with concrete outputs and gates.

Agentic Development Lifecycle (ADLC)



Agentic Development Lifecycle (ADLC)

Stage 1: Ideate and Evaluate

The purpose of Stage 1 is to define *what* the agent is for, *where* it is allowed to act, and *how* success and failure will be measured—before any agent composition begins.

In SDLC, teams start with features and requirements. In ADLC, begin with ideation and intent: a North Star outcome, bounded by policy and risk tolerance. The goal is to prevent the most common agent failure pattern: building impressive demos that don't change workflows or survive production constraints.

In the Ideate and Evaluate stage, teams should begin by defining North Star goals that express desired business outcomes (KPIs) rather than technical capabilities. This ensures that the agent is built to change workflows, not just demonstrate clever functionality. Alongside these goals, teams should articulate behavioral questions that describe how the agent should act in moments of ambiguity, conflict, or incomplete information—because these edge cases are where trust is won or lost.

Equally important is a risk and failure-mode map that anticipates what could go wrong, how problems will be detected, and which conditions should trigger a “stop-the-line” response or human escalation. Finally, an explicit evaluation plan defines what “good behavior” looks like in practice, how it will be tested through scenarios and drills, and which metrics actually matter to the business. Taken together, these outputs reflect an evaluation-first approach: teams define success, risk, and measurement upfront, then design the agent to meet those standards—rather than retrofitting governance and evaluation after the system is already in motion.

As laid out in this book, agentic transformation is an operating-model shift: reimagining workflows as AI-first, with humans selectively reintroduced for judgment, policy, and accountability. That means your Ideate and Evaluate stage should explicitly define where humans must remain accountable and what escalation thresholds look like.²⁹

Stage 2: Compose (Build)

In this stage, composing an agent is like assembling a disciplined team—consisting of models + instructions + tools + memory + guardrails—rather than “coding” a single deterministic program.

- The following composition decisions create most of your downstream success or failure:
- **Model selection:** capability, latency, cost, privacy posture, jurisdictional constraints
- **Persona and instruction design:** system prompts, policy prompts, style, refusal rules
- **Tool integration:** APIs, databases, search, workflow triggers—each with scoped permissions
- **Memory design:** what has persisted, where, for how long, and under what privacy/records rules
- **Guardrails by design:** allowed actions, forbidden actions, data boundaries, and audit requirements

What changes most from the traditional SDLC is the very nature of what teams are building and controlling. It becomes a composition of prompts, policies, tools, retrieval mechanisms, orchestration layers, and runtime controls that collectively shape how the agent reasons and acts. Because behavior emerges from the interaction of these components, even small changes—such as an untested prompt adjustment or a new tool integration—can have a disproportionately large blast radius, affecting volumes of interactions in production. This is why agent delivery requires SDLC-like discipline adapted for probabilistic systems, with stronger emphasis on evaluation, controlled change management, and staged promotion across environments.

To manage this complexity, a formal composition gate is essential before agents are promoted beyond development and testing. All tools and integrations should be registered in an approved catalog, with permissions defined according to least-privilege principles. Data sources must be clearly classified so that agents can distinguish between authoritative systems of record and untrusted or advisory sources. Finally, teams should define a traceability plan that specifies what interactions, decisions, and actions will be logged and retained, ensuring that agent behavior can be audited, explained, and governed over time.

Stage 3: Test (Behavioral Drills)

In the Test stage, we shift from bug-finding to behavioral validation—stress-testing reasoning, tool use, safety, and goal alignment.

For agents, testing means continuous validation under real-world conditions, including edge cases and compliance checks.³⁰ In this stage, sandboxing and simulated environments can be used to observe decision-making before making any real-world impact.

Core test types include:

1. **Scenario drills (business realism):** messy inputs, conflicting policies, incomplete records.
2. **Tool-use drills:** wrong-tool temptation, tool failures, permission denials, rate limits.
3. **Grounding and hallucination tests:** does the agent invent facts, cite the wrong source, or overclaim?
4. **Alignment and policy tests:** does it respect boundaries (privacy, records, approvals)?
5. **Adversarial drills:** prompt injection, malicious instructions embedded in retrieved content, data poisoning attempts.

Agent systems require evaluation metrics that go well beyond traditional operational measures such as uptime and latency. Useful metrics include context relevance, faithfulness, and answer similarity—measures that assess whether an agent is grounding its responses in the right sources, staying faithful to the available evidence, and producing outputs that are consistent and reliable. Even in deployments that are not explicitly retrieval-augmented (RAG-first), these metrics generalize well because they capture the underlying quality of agent behavior: whether the system is using appropriate context, avoiding hallucination, and delivering trustworthy outcomes.

To ensure these standards are met before agents are promoted into broader operational use, teams should establish a formal test gate. This gate requires that the behavioral drill suite passes across representative scenarios, including normal, edge, and stress conditions, and that red-team findings (from testing for unsafe behavior) are documented along with clear mitigations. In addition, any expansion of agent access or autonomy should be explicitly approved, since agents typically begin with constrained tool scope and permissions that are only broadened once they have demonstrated stable, trustworthy behavior.

Stage 4: Deploy (Staged Rollout)

Deployment in the agentic lifecycle should be treated less like a traditional software release and more like onboarding a new employee on their first day of work. The goal is not to “fire and forget,” but to introduce agents into live environments gradually, with supervision, clear boundaries, and the expectation that they will need guidance, feedback, and time to prove reliability.

Just as a new hire is not given full authority on day one, agents should begin with limited scope and carefully defined permissions, operating under close observation. This staged onboarding approach acknowledges that agent behavior will evolve and that early production exposure is part of the learning process—but only when paired with strong guardrails, human oversight, and the ability to intervene quickly if behavior deviates from expectations.

A staged rollout pattern includes:

1. **Development:** agent composition and early drills with synthetic or masked data.
2. **Staging:** production-like integrations + test data + load and failure simulations.
3. **Production (limited):** constrained permissions, limited user cohort, tight monitoring, human-in-command gates.
1. **Production (expanded):** permissions and scope expand only as behavioral performance proves stable.

Throughout this progression, formal deployment gates should be enforced, including explicit approval for promotion into production, runtime monitoring that captures both behavioral and operational signals, and clearly defined emergency stop or shutdown mechanisms to contain risk if an agent behaves unexpectedly in live environments.

Stage 5: Observe and Optimize (Continuous Monitoring)

Stage 5 focuses on continuously monitoring agent behavior in production to detect behavioral drift, tool misuse, policy breaches, and performance against key business outcomes, and then feeding those insights back into earlier stages of the lifecycle.

Post-deployment observation is not limited to technical uptime or throughput; it encompasses how agents are actually being adopted, how users interact with them, and how they access and act on sensitive data. Sustained value depends on monitoring performance, adoption patterns, user behavior, and data access after launch so teams can optimize continuously rather than treating deployment as the end of the journey. This is necessary because agentic systems introduce heightened risks around traceability, bias, hallucinations, and security—risks that become more pronounced as agents operate with greater autonomy and less direct supervision.

Effective observation requires tracking both outcome and behavioral signals. Outcome KPIs such as cycle time reduction, quality improvement, throughput, cost-to-serve, and compliance rates indicate whether agents are delivering tangible business value. Behavioral KPIs—including escalation frequency, appropriateness of refusals, quality of tool selection, and early indicators of behavioral drift—help teams understand how agents are making decisions in practice.

Risk signals such as anomalous tool usage, unusual data access patterns, or repeated near-miss behaviors provide early warnings of emerging issues, while monitoring human-in-command load reveals whether people are primarily handling exceptions or are being pulled back into routine work the agent was meant to absorb.

Crucially, observation is not a passive reporting activity but the engine of iteration: insights from production feed back into Stage 1 as new behavioral questions; into Stage 2 as re-composition of models, tools, or policies; and into Stage 3 as new drills and evaluations, ensuring the system evolves in step with real-world conditions.

The Agent's Experimentation Loop

This experimentation loop typically cycles through variations in models, prompts and policies, tool configurations, retrieval strategies, and guardrails, with each iteration designed to test a specific behavioral hypothesis: how the agent reasons under ambiguity, how reliably it selects the right tools, how it handles edge cases, and how well it adheres to defined policies and safety constraints. The goal is not to perfect the agent in one pass, but to progressively shape its behavior through disciplined experimentation before it ever encounters real users or live systems.

The experimentation loop depends on strong observability and feedback. Each build-and-test cycle should generate structured evaluation data, including behavioral drill results, failure patterns, and alignment metrics, so teams can compare versions, understand regressions, and make evidence-based decisions about what to change next. Over time, this loop becomes a learning system for the organization as much as for the agent itself, capturing institutional knowledge about which compositions, prompts, tools, and guardrails produce stable, trustworthy behavior. When this experimentation loop is weak or informal, teams tend to develop fragile agents whose behavior degrades in production; when it is rigorous and instrumented, it becomes the foundation for safe autonomy at scale.

The Agent's Runtime Optimization Loop

While the ADLC governs how agents are designed, tested, deployed, and improved over time, each agent also operates through a continuous internal loop during execution. In practice, agents move through a repeating cycle of perceiving signals from users, enterprise content, telemetry, and external data sources; reasoning about that context and forming a plan by breaking goals into actionable sub-tasks; acting through tools, APIs, and workflows to carry out those plans; and then reflecting on outcomes to evaluate success, log evidence, and update memory or policies for future decisions. This loop is not a one-time sequence but the core operating rhythm of agentic systems in production.

This internal loop is central to how autonomy is governed in practice. Monitoring and governance mechanisms must be designed around these steps, providing visibility into what the agent perceived, how it reasoned, which tools it selected, and what outcomes those actions produced. If teams cannot observe and audit tool selection and action results, they effectively lose the ability to govern autonomous behavior, diagnose failures, or enforce accountability. In agentic systems, observability of the loop is what makes autonomy manageable rather than opaque.

Governance and Security

Agentic AI governance is not a checklist at the end. It is a continuous set of controls across *design, build, test, deploy, and run*.

Governance primitives that must exist end-to-end include:

- **Verifiable agent identities** (so actions can be attributed)
- **Least-privilege authorization boundaries** (per tool, per environment, per task class)
- **Audit trails and lineage** (who changed what, which version ran, what actions occurred)
- **Policy-as-code** (for approvals, thresholds, and data access)
- **Evaluation telemetry** (behavioral and grounding metrics at runtime)

Governance must become real-time and embedded, with humans maintaining final accountability while shifting from line-by-line review to policy definition, outlier monitoring, and calibration of human involvement.

The transition from theoretical ADLC stages to industrial-scale execution is exemplified by the following case study of a Middle Eastern customs agency, which moves beyond traditional linear deployment to an evaluation-first ecosystem where every border-flow automation is treated as a continuous behavioral loop rather than a static software release.



A Middle Eastern Customs Agency



By embedding continuous evaluation and feedback into our delivery lifecycle, we shifted from reactive testing to proactive quality assurance, allowing our teams to catch issues early, adapt quickly, and deliver more reliable services with greater confidence.

– Senior Technology Lead, Customers

A pivotal Middle Eastern government agency oversees trade and border flows in one of the world's busiest logistics hubs. Under mounting pressure to maintain high-visibility digital services while upholding security and compliance, the organization moved beyond fragmented, manual testing processes toward an integrated lifecycle where outcomes, automation, and continuous evaluation drive behavioral quality. The agency implemented an AI-powered functional testing lab to create a unified delivery platform that supports iterative feedback loops and real-time insight into quality across the entire delivery pipeline.

This evolution mirrors the core stages of ADLC: ideation and outcome definition focused on reliable delivery velocity, composition of integrated tooling and test automation, rigorous behavioral drills across scenarios, staged deployment with governance gates, and ongoing observation to detect drift or regressions. By anchoring delivery around measurable business outcomes and continuous validation rather than isolated task completion, teams at the customs agency achieved faster release cycles and stronger confidence in production outcomes. The result is a delivery ecosystem that blends automation, traceability, and adaptive learning—positioning the organization to scale digital services securely and responsively in a rapidly changing operational environment.

The agency's success demonstrates that in the high-stakes world of international logistics, the ADLC serves as the vital bridge between speed and security. By institutionalizing the "Experimentation" and "Runtime Optimization" loops, the agency didn't just automate existing manual processes; they created a self-correcting infrastructure capable of evolving alongside global trade volatility. This case serves as a blueprint for any sovereign entity looking to deploy autonomous agents: when governance is embedded into the lifecycle rather than treated as a final hurdle, the result is a resilient, always-on delivery ecosystem that scales without compromising public trust.

ADLC Challenges

Organizations face a distinct set of challenges when operationalizing the Agentic Development Lifecycle. One of the most common is scope distortion driven by hype—where teams apply agentic approaches to problems that do not require autonomy or attempt to overlay intelligent agents onto workflows that are poorly designed or structurally unsound. This misalignment between the use case and the level of capability deployed is a leading cause of stalled, retrenched, or canceled agentic AI initiatives.

Data readiness and policy foundations present another significant barrier. Responsible autonomy depends on structured, traceable, and interoperable data, along with clear policy constructs; without these foundations, agentic systems tend to amplify trust, compliance, and regulatory risks rather than create sustainable value.³²

Operational risks also intensify as agents move into production. Behavioral drift and runaway variance can emerge over time as environments change, tools evolve, or underlying models, prompts, and data sources shift, which is why continuous observation and evaluation are essential to maintaining stable behavior. Once agents are empowered to act in real systems, misconfigured integrations or flawed decision-making can quickly escalate into operational incidents, including exposure of sensitive data or irreversible actions.

Finally, governance becomes increasingly complex at scale. As organizations deploy heterogeneous estates of agents across multiple vendors, models, and business domains, point governance approaches break down. An agent control plane highlights the need for consistent, out-of-band oversight to maintain visibility, enforce policy, and manage risk across the entire agent landscape. Unlike a development environment or orchestration layer, this plane provides cross-cutting oversight across a heterogeneous, multi-vendor agent landscape, ensuring autonomous behavior remains aligned with business intent, policy constraints, and risk tolerance.³³

ADLC Best Practices.

- **Start with “evaluation-first” outcomes and failure modes.** Write the behavioral questions, success metrics, and unacceptable failure modes before composing the agent.
- **Treat environments as non-negotiable.** Never test in production. Use dev/staging/prod separation with explicit promotion gates and approval workflows.
- **Use behavioral drills as your regression suite.** Maintain a living library of scenarios that represent your real operating conditions—and rerun them after every prompt/tool/model change.
- **Constrain autonomy early; expand deliberately.** Launch “rookie agents” with limited permissions and human-in-the-loop triggers. Expand scope only after stable performance.³⁴
- **Choreograph the agent’s operational loop.** Log perception inputs (appropriately), plans, tool calls, outcomes, and escalation events—so you can explain behavior and debug failures.
- **Adopt out-of-band oversight for multi-agent estates.** Plan for an agent control plane that provides consistent visibility, policy enforcement, and risk management across the estate.
- **Make governance a growth enabler, not a brake.** Governed autonomy is what creates confidence to scale.
- **Pick the right use cases and redesign workflows.** Pursue agentic AI only where it delivers clear ROI. This requires rethinking workflows from the ground up rather than forcing agents into legacy process patterns.

Executive Checklist: Are You Ready to Operationalize ADLC?

Use this checklist to assess whether your organization is ready to design, deploy, and govern agentic AI systems at enterprise scale.

1) Strategy and Intent (Ideate and Evaluate)

- Have we defined a clear **North Star business outcome** (not just an AI capability)?
- Have we articulated the agent’s **intent, scope, and boundaries**?
- Have we documented **behavioral questions** (how the agent should act under ambiguity, pressure, or conflict)?
- Do we have **success metrics** tied to business KPIs (e.g., cycle time, quality, compliance), not just model accuracy?
- Have we identified **high-risk failure modes** and “stop-the-line” conditions?

2) Composition and Integration (Build)

- Are foundation models selected based on **capability, cost, latency, privacy, and jurisdictional constraints**?
- Are agent personas, system prompts, and policies **versioned and reviewable**?
- Are all tools/APIs registered in an **approved tool catalog** with least-privilege access?
- Is **agent memory clearly defined** (what is stored, where, for how long, and under what records/privacy rules)?
- Do we have guardrails for **what the agent is explicitly not allowed to do**?

3) Testing and Evaluation (Behavioral Drills)

- Do we have a repeatable suite of **behavioral drills** covering normal, edge, and adversarial scenarios?
- Are we testing **tool-use behavior**, refusal behavior, and escalation behavior—not just output quality?
- Do we measure **grounding and reliability** (e.g., relevance, faithfulness, hallucination rates)?
- Are **red-team results** documented with clear mitigations?
- Is there a formal **promotion gate** from test to production?

4) Deployment and Operations (Staged Rollout)

- Are agents deployed through **segregated environments** (dev → staging → production)?
- Are human-in-command controls defined for **high-risk decisions**?
- Are **permissions constrained at launch** and expanded only after stable performance?
- Do we have a **kill switch / shutdown mechanism** for emergencies?
- Is **operational ownership** clearly assigned (who is on the hook when agents fail)?

5) Observation, Learning, and Optimization (Continuous Monitoring)

- Are we monitoring for **behavioral drift**, not just uptime and latency?
- Do we track **business impact KPIs** alongside technical metrics?
- Are agent actions and tool calls **logged and auditable**?
- Do we have a formal feedback loop from production back into **Ideate & Evaluate**?
- Are we periodically **reviewing whether humans are still doing work** the agent should be doing—or vice versa?

6) Governance and Security (Applies to Every Stage)

- Does each agent have a **verifiable identity**?
- Are authorization policies defined and enforced at the **tool and data layer**?
- Are prompts, configurations, and policy changes **version-controlled**?
- Can we reconstruct **who did what, when, and why** for audits and investigations?
- Is governance treated as a **scaling enabler**, not a deployment blocker?

Red Flags (If You Answer "Yes" to Any of These, Pause Before Scaling)

- We cannot explain **why** the agent behaved the way it did in production.
- The agent has direct access to critical systems **without staged rollout or approvals**.
- Success is **measured only by "the demo worked,"** not by sustained business outcomes.
- Governance was **added after deployment** rather than designed in from day one.

Executive Implications

CAIO

Safe agent deployment requires governance across the full lifecycle. The CAIO defines the risk classification, policy oversight, and audit controls that determine when agents can enter, operate within, and exit production.

CIO

Agentic AI must be managed as a living platform—versioned, monitored, and governed continuously—rather than as static software releases.

CFO

The total cost of agent ownership includes training, monitoring, drift remediation, and retirement; lifecycle discipline is required to prevent AI portfolios from becoming a hidden operating expense.

CHRO

As agents mature and assume operational responsibilities, workforce models must evolve to include supervision, exception handling, and stewardship of digital labor.

CDO

Model drift and data decay directly undermine agent performance; continuous evaluation and lineage-aware retraining are required to preserve trust and decision quality over time.

COO

Operational reliability depends on managing agent behavior over time, with clear escalation, rollback, and retirement mechanisms to prevent automation debt from accumulating.

In the next section of this book, Part 4, we'll discuss the execution and adoption of AI, including a phased adoption roadmap and how to measure your return on agentic AI investment.

Part 4

Execution and Adoption

04


Chapter Eleven

From Pilots to Platforms

In this chapter, we address the sobering reality of the enterprise AI boom: most GenAI initiatives have failed to meet their projected returns. Drawing on research from top analysts, we examine why fragmented deployments inevitably stall. More importantly, we show how the architectural, organizational, and governance frameworks developed throughout this book provide a practical pathway from costly experimentation to durable operational value.

The ROI Reckoning

The enterprise AI experiment has reached a sobering verdict. Despite global AI spending rising to \$1.5 trillion in 2025, financial returns remain underwhelming. Only 17% of organizations attribute more than 5% of EBIT to GenAI, and just 1% consider their strategies mature.³⁵ As stated in Chapter 1, 95% of enterprise GenAI pilots fail to produce measurable P&L impact.³⁶



The GenAI Divide: AI Leaders vs. AI Laggards³⁷

On average, enterprises abandon 46% of AI proofs of concept before they ever reach production. Among large organizations, only 25% of AI initiatives deliver the ROI leaders expected, and just 16% successfully scale across the enterprise.³⁸ For the executives who have championed these digital transformations, these figures require honest accounting. This widespread stalling is not happening because the underlying technology is flawed. It's happening because the initial deployment paradigm was structurally incapable of delivering the promised value. Understanding this architectural mismatch is the absolute prerequisite to getting agentic AI right.

Why Pilots Stall


The gap between AI ambition and outcomes is not a mystery. Research converges on a consistent set of failure patterns. What is striking is how few of these are fundamentally technical problems. The vast majority are organizational, architectural, and governance failures—precisely the domains that are addressed in this book.

Let's drill down on the main reasons why AI pilots fail:

- 1. Pilot Purgatory.** Roughly 90% of vertical, function-specific GenAI use cases remain stuck in pilot mode.³⁹ Fewer than one-third of GenAI experiments have moved into production, and large enterprises take nine months on average to move from pilot to production. The pattern is consistent: impressive demos that never graduate to operating value.
- 2. Data as the Ultimate Governor.** A staggering 85 percent of AI project failures are directly attributed to poor data access and quality.⁴⁰ At the GenAI layer, data management is already a leading constraint; in agentic architectures, it becomes existential. Data is the cognitive substrate on which autonomous digital actors are trained and the primary context they use to make decisions. Training agents on ungoverned, siloed, or low-quality data does not create intelligence—it automates and amplifies organizational blind spots. When agents act at machine speed, data integrity determines the safety and defensibility of the enterprise. As argued in Chapter 2 of this book, information is the fuel of the cognitive operating model. When that fuel is contaminated, the engine doesn't stall—it accelerates the organization toward systemic, machine-scale failure.

For GenAI high performers, data management is the top challenge in capturing value from the technology.

Elements that have posed challenges in capturing value from generative AI, % of respondents




Data As Top Challenge in GenAI⁴¹

- 1. Cost Escalation Without Returns.** Seventy-three percent of companies spend at least \$1 million annually on GenAI. However, due to a lack of ROI, CFO confidence has eroded, with fewer than 27% planning to increase GenAI budgets, down from 53% the year before.⁴² When costs rise and outcomes do not, the investment thesis collapses.
- 2. The Chatbot Ceiling.** The core limitation of first-generation GenAI is architectural. Copilots (or AI-powered assistants) and chatbots operate in a prompt-and-response paradigm: they cannot take independent action, integrate deeply with backend systems, retain organizational memory, or produce the deterministic outcomes enterprises require. As Chapter 4 established, accelerating work at the edge cannot overcome friction in the core. Making individuals write emails or code faster does not redesign end-to-end workflows—or change enterprise performance.

- 1. Integration Fragmentation.** Early GenAI models were deployed without redesigning the workflows they were meant to improve, leading to a proliferation of disconnected pilots. While each showed promise in isolation, together they failed to produce coordinated, systemic outcomes. Intelligence was bolted onto the interface, not embedded in the operating model.
- 2. Governance as an Afterthought.** With only 18% of organizations reporting an enterprise-wide responsible council, governance is still being treated as a downstream compliance exercise rather than an architectural prerequisite for scale.⁴³ As Chapter 3 posits, defensibility is not an afterthought for agentic AI—it is a core design requirement. Organizations that attempt to bolt governance onto autonomous deployments after the fact will inevitably discover a hard truth: trust cannot be retrofitted.
- 3. Organizational Misalignment Between Technology and Business Leadership.** We've discussed the structural imbalance in earlier chapters: technology teams own 53% of enterprise AI efforts, while business functions own just 7%. CIOs and CTOs hold 44% of AI technology strategy responsibility but only 29% of AI business strategy.⁴⁴ This concentration explains why pilots succeed technically but fail operationally—there is no business ownership of outcomes. When AI remains perceived as “an IT program,” business adoption, accountability, and operating model transformation stall. The data shows that cross-functional alignment lifts portfolio success probability from 76% to 85%—a nine-point improvement representing the single highest-leverage decision available to enterprise leadership. Organizations that fail to distribute AI ownership across technology and business functions risk concentrating capability in a small expert team while demand for AI-driven outcomes outpaces delivery capacity.

Nearly one-quarter of respondents say their organizations have experienced negative consequences from generative AI's inaccuracy.

GenAI-related risks that caused negative consequences for organizations, % of respondents



Asked only of respondents using GenAI in ≥1 function (n=876). 'Don't know/not applicable' (17%) not shown.
 Source: McKinsey Global Survey on AI, 1,363 respondents, Feb 22–Mar 5, 2024.

The Structural Mismatch

While 78% of companies have deployed generative AI in some form, more than 80% of these report no material earnings impact.⁴⁶ The gap is structural, not technical. Horizontal copilots scaled quickly, but their gains were diffuse, difficult to measure, and spread thinly across individuals. Meanwhile, the high-value, vertical use cases that require deep integration and workflow redesign remained trapped in pilot mode.

As outlined in Chapter 5, enterprise value is created through the coordinated flow of information, decisions, and actions across end-to-end processes. A procurement agent may summarize a vendor contract in seconds, but if onboarding still requires manual data entry across disconnected compliance systems, the supply chain remains bottlenecked.

In practice, GenAI has often been deployed in isolated sandboxes, unable to act through backend systems, trigger workflows, or retain organizational memory. Each interaction started from zero. These are not limitations solved by better prompts; they are architectural constraints inherent to the copilot paradigm.

The resulting disillusionment is not a failure of AI itself, but of the first deployment model. The value enterprises expected from early AI investments runs through a fundamentally different architecture.

From Copilots to Agents: The Architectural Pivot

Agentic AI is one of Gartner's Top Strategic Technology Trends for 2025:

Many startups are already marketing themselves as AI-agent-building platforms, some including AI agents that help build users' AI agents from natural language workflows written by line-of-business owners. Hyperscalers are adding agentic AI to their AI assistants.⁴⁷

Agentic AI is not an incremental improvement on generative AI; it is a paradigm shift. Where early GenAI produced outputs in response to prompts, agentic systems plan, reason, decide, and act to complete complex, multi-step goals. This shift has rapidly reshaped enterprise architectures, with major platforms reorienting around agents as the new execution layer.

This evolution moved through three distinct stages. The first was the AI assistant or chatbot era of 2022 to 2023, characterized by reactive, single turn, isolated interactions. The second was the copilot era of 2023 to 2024, where AI was embedded in workflows to suggest actions, but remained entirely human dependent at every step. The third is the agentic era emerging in 2025 and beyond: proactive, goal driven, deeply integrated, and capable of autonomous execution across enterprise systems and silos.

Agents possess execution capabilities that chatbots lack. They take direct action to execute end-to-end workflows, orchestrate multiple agents on complex objectives, and continuously learn to adjust strategies based on outcomes. This is the architecture developed throughout this book: nested orchestration (Chapter 3), the Agentic Genome Map (Chapter 5), reference architecture (Chapter 8), deployment models (Chapter 9), and the agent lifecycle (Chapter 10). Each element addresses the structural failures that limited early GenAI. Agentic AI is not an iteration of a GenAI model—it is a different *operating* model altogether.

Why Fragmented AI Fails at Scale

The pattern is now well-documented. Organizations that deploy AI as point solutions—isolated agents for individual tasks without workflow redesign or shared infrastructure—repeat the same failure mode that limited GenAI, only at greater cost and operational risk. As Chapter 9 explains, pilots can succeed without defensibility; operating models cannot. When agents operate in silos, predictable failures follow:


- **Duplication and complexity:** Teams build narrowly scoped agents for individual tasks using multiple platforms, resulting in overlapping capabilities, inconsistent behaviors, and unnecessary costs. Without a unified platform approach, the agent estate quickly becomes unmanageable.
- **Governance gaps widen:** Point solutions can't be governed consistently. When each agent operates under its own rules with independent data access and isolated logging practices, the enterprise loses the ability to prove what happened and why. In regulated industries, this exposure becomes existential.
- **Integration remains brittle:** Isolated agents cannot coordinate end-to-end workflows. Without orchestration binding intelligence to execution, the same handoff friction that constrained copilots persists. Accelerating a single task is meaningless if core processes remain fragmented and bottlenecked by manual handoffs.
- **Learning loops fail to form:** Without feedback mechanisms connecting production behavior back to design and testing, agents stagnate while the business evolves around them. As described in Chapter 10, continuous experimentation and runtime optimization loops are what transform agents from static tools into adaptive systems.
- **Trust erodes:** Users quickly disengage when outputs feel inconsistent or difficult to verify. Over time, this erosion of trust negates any efficiency gains achieved through automation. In contrast, when AI outputs are reliable, high-performing teams treat agents less like software features and more like new employees, assigning them defined roles, clear performance expectations, and structured evaluation frameworks.

Evidence from recent enterprise AI deployments shows that organizations pursuing agentic architectures in isolation face materially higher failure rates than those leveraging platform partnerships. Three out of four firms attempting to build advanced agentic architecture will fail. Vendor partnerships succeed approximately 67% of the time, compared to only 33% for fully internal builds.⁴⁸ Fragmented approaches to agentic AI do not compound value; they compound risk.

From Pilots to Platforms: The Operating Model Shift

Moving from isolated pilots to enterprise platforms is not a tooling upgrade; it is an operating model transformation. The consistent lesson across early deployments is that durable value emerges only when intelligence is embedded into workflows, grounded in governed enterprise memory, orchestrated as a shared platform, and scaled with built-in trust and organizational change.

Seventy percent of AI success depends on people, processes, and organizational culture. Organizations that invest deeply in change management achieve 60% adoption rates, compared to just 30% for those that do not.⁴⁹ This requires dedicated adoption teams, role-based training, senior leaders actively modeling AI use, phased rollout roadmaps, and exceptionally clear accountability structures.



The 10-20-70 Model: 70% of AI Success Depends on People, Processes, and Organizational Culture⁵⁰

Here are five basic platform principles:

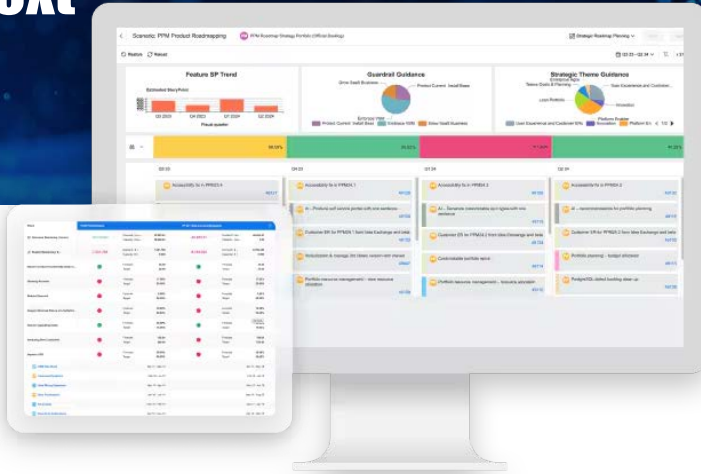
- 1. Lead with workflows (not tools).** Anchor the transformation in high-value, end-to-end workflows (e.g., order-to-cash, procure-to-pay, incident-to-resolution), not isolated use cases or models. This aligns with the nested orchestration model (Chapter 4) and the Agentic Genome Map (Chapter 5), while anchoring the organizational model required for scale.
- 2. Build the enterprise memory first.** Agents are only as capable as the enterprise memory they are allowed to reason over. Data quality remains the single most cited barrier to AI success across every major survey. As established in Chapter 8, content platforms must be explicitly designed as services, rather than isolated silos, with rich metadata, clear lineage, and strict governance in place before agents act. Without this foundation, you are building an intelligent system on sand.



- 1. Build orchestration as a shared platform.** Orchestration converts raw intelligence into measurable outcomes by binding agents to execution across systems of record. Through orchestration, metrics such as latency, throughput, cost-to-serve, SLA adherence, and exception rates become observable properties of intelligent workflows.
- 2. Govern autonomy by design.** Governance is not a compliance phase to be bolted on after deployment. It is the continuous control plane spanning design, build, testing, deployment, and runtime. Chapter 10's Agentic Development Lifecycle makes this explicit. Organizations that treat governance as a scaling enabler, rather than a deployment blocker, build the trust required to safely expand autonomy.
- 3. Scale through agent portfolios, not one-offs.** Mature AI programs invest in reusable agent capabilities and shared services, such as document ingestion, classification, retrieval, validation, and compliance checks, which can be dynamically recomposed across multiple workflows. Adopt a strict registry mindset consisting of approved agents, approved tools, approved workflows, and measured outcomes.
- 4. Establish cross-functional AI coordination from day one.** Data shows AI governance requires active collaboration across the CAIO, CIO, CISO, CTO, CHRO, CDO, and business stakeholders. The five principles above address architecture, data, orchestration, governance, and portfolio management—but without an organizational coordination mechanism that connects Engineering, Sales, Marketing, HR, Compliance, and Professional Services around a single vision, these capabilities risk being executed in fragments. Whether through a dedicated AI strategy function, an autonomous steering committee, or another mechanism that leadership deems appropriate, this coordination is the prerequisite for everything else described in this chapter.

While 95% of enterprise GenAI pilots fail to move the needle on the P&L, OpenText provides a roadmap for the remaining 5% by demonstrating how a unified platform architecture can collapse 180 fragmented initiatives into a single, high-velocity orchestration engine that delivers immediate operational savings.

OpenText



When a global team at OpenText faced fragmented project tracking across more than 180 active initiatives, disparate tools like Jira, Excel, and Microsoft Project made it difficult to achieve consolidated visibility, coordinate work, or deliver consistent outcomes. To address these barriers, the Cloud Services team deployed OpenText Project and Portfolio Management (PPM) as a unified platform to connect strategy with execution and standardize project tracking across engineering and operations.

Rather than managing project status and workstreams in isolated systems, the PPM solution automated previously manual reporting tasks, generating consolidated status updates and delivering an enterprise-wide view of execution progress. This shift enabled teams to eliminate redundant effort—saving an estimated 360 hours per week—and laid the foundation for AI-driven project automation and intelligent work orchestration supported by a shared content and portfolio substrate.

By establishing a single source of truth for work execution, the organization effectively replaced disconnected task management with a coordinated workflow platform that aligns planning, monitoring, and automated reporting. This not only improved operational visibility and execution discipline—but also set the stage for future capabilities where intelligent automation agents can participate in cross-project workflows, monitor performance signals, trigger actions, and enforce governance across the enterprise portfolio.



Workforce Capacity Reclaimed: 360 hours per week saved.



Scale of Coordination: 180+ active initiatives unified under a single orchestration layer.




Administrative Efficiency: Transitioned from manual reporting to a zero-touch, real-time reporting environment.

By solving for integration fragmentation and establishing a single source of truth today, OpenText has effectively future-proofed its operating model—ensuring that as they transition from copilots to fully autonomous agents, those agents will inherit a governed, high-quality data substrate rather than becoming yet another isolated pilot lost in purgatory.

Early Evidence: Agentic AI Delivers Measurable Returns

Early indicators show AI agents delivering 13.7% ROI—outperforming non-agentic GenAI and clearing typical enterprise cost-of-capital thresholds. At scale, agents are projected to drive 3–5% annual productivity gains, with this baseline accelerating as the duration and complexity of tasks they reliably execute continues to double every four months.⁵¹

2x more value is expected to come from agentic AI by 2028



Companies are experimenting with agentic AI...

46%

Of companies are experimenting with early pilots or deploying agents
(with 16% of these companies already demonstrating tangible value)

...Thanks to budget allocation

30%

Of companies are spending over 15% of their AI budget on agents

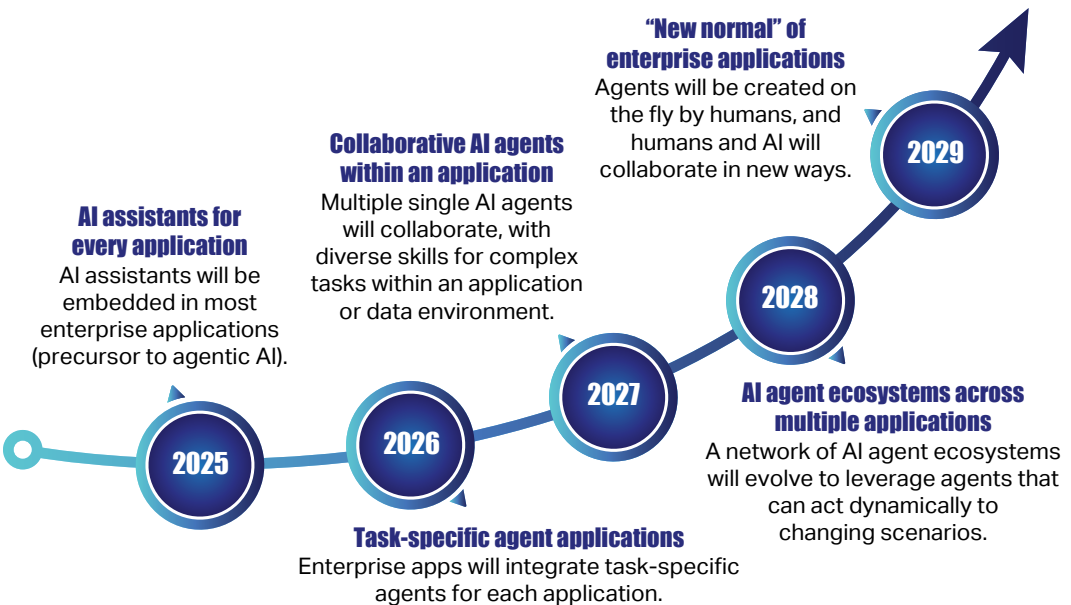
Value from Agentic AI Is Expected to Double by 2028⁵²

Case studies across industries are demonstrating concrete, measurable returns:

- A shipbuilder reduced engineering efforts by 40% and cut design lead times by 60% using multi-step agentic processes.
- A telecom operator deployed agentic assistants that produced a 5× increase in digital sales.
- An industrial company doubled its prospecting efforts, driving a 40% increase in order intake.
- Automation in insurance claims processing achieved 245% ROI, and financial services automation delivered over 250% ROI.⁵³

Market projections reflect this momentum. Gartner projects that AI agents will intermediate \$15 trillion in B2B purchases by 2028—commerce that flows through trading networks requiring governed transaction integrity.⁵⁴ Simultaneously, 40% of agent projects will be canceled by 2027 due to 20–30x token cost overruns, creating CFO urgency for agent governance and FinOps capabilities.⁵⁵ The context-as-a-service market alone—governed content delivery for agentic consumption—represents a \$52 billion addressable market by 2030. The adoption of agentic AI is outpacing every prior technological wave.

These results come with a critical caveat. Organizations realizing measurable returns are not just deploying better models, they are redesigning workflows, governing their data substrate, orchestrating intelligence across systems, and treating deployment as an operating model transformation. Technology is necessary, but not sufficient. Architecture, governance, and organizational readiness ultimately determine whether enterprise value is realized.



The Future of Agentic AI in Enterprise Applications⁵⁶

The Compounding Cost of Delay

Unlike prior technology waves, agentic AI creates compounding advantages that widen over time. Each interaction generates proprietary learning, and workflow redesign builds organizational capabilities that take years to replicate. Talent dynamics amplify this lead: with a 3:1 demand-supply gap for AI professionals, pioneers attract scarce expertise while laggards face rising costs.

The emerging pattern is not incremental improvement, but structural divergence. Only 5% of companies are considered “future-built” and generating substantial AI value, while 60% remain laggards with minimal gains. Future-built firms achieve 1.7× revenue growth and 3.6× three-year total shareholder returns. They allocate significantly more of their AI budgets to agents and actively deploy them, compared to zero among the laggard group.⁵⁷

For the public sector, the stakes are equally high. AI can free up 30% of public servants’ time, delivering a triple dividend—lower fiscal deficits, higher administrative productivity, and GDP growth. The risk of inaction is a widening capability gap between public and private sectors that will be difficult to reverse.⁵⁸

The strategic imperative is clear. The window for building the architectural foundations, organizational capabilities, and governance frameworks required for agentic AI at scale is currently open, but it is narrowing rapidly. Organizations that delay face not merely a late start, but a rapidly widening capability gap. For enterprises operating in the European Union, the window is not just narrowing competitively—it is closing regulatorily. Under the EU AI Act, organizations deploying AI in regulated domains will soon be required to demonstrate documented risk controls, oversight mechanisms, and complete operational traceability—making the potential penalty of up to 7% of global revenue the most concrete cost-of-delay data point for enterprise leadership today.

What This Book Has Built: Architecture as the Answer

This book defines a coherent architectural response to the failures that limited GenAI’s enterprise value. Each element directly addresses a root cause of the ROI gap:

1. **The pilot purgatory problem** is dismantled by the nested orchestration model in Chapter 4 and the Agentic Genome Map in Chapter 5. These provide a structural framework for scaling intelligence from individual atomic tasks to end-to-end workflows. Intelligence designed for orchestration from the outset doesn’t get trapped in isolated demonstrations.

2. **The data quality problem** is resolved by the content substrate architecture detailed in Chapters 2 and 8. This establishes governed enterprise memory as the absolute foundation for agent reasoning. Agents grounded in authoritative, metadata-rich content platforms produce highly defensible outcomes. Conversely, agents operating on ungoverned data produce unacceptable risk.
3. **The integration problem** is addressed by the orchestration layer in Chapters 4, 5, and 8. This layer forcefully binds intelligence to execution across systems of record. Orchestration systematically converts individual agent capabilities into coordinated, easily measurable workflows.
4. **The governance problem** is solved by the sovereignty and trust frameworks covered in Chapters 3, 8, and 9, alongside the agent lifecycle in Chapter 10. These embed strict governance into every single stage of agent design, composition, testing, deployment, and operation. Trust becomes a measurable architectural property, not a vague aspiration.
5. **The organizational readiness problem** is managed by the human-in-command design principles in Chapters 3, 4, and 7, and the deployment models in Chapter 9. These deliberately structure human-agent collaboration, strict escalation pathways, and clear accountability frameworks.
6. **The measurement problem** is tackled by the evaluation-first approach of the ADLC and the executive scorecard in Chapter 10. These aggressively anchor agent performance to tangible business outcomes rather than isolated model metrics. What gets measured gets scaled.
7. **The cross-functional alignment problem** is addressed by the organizational design principles in Chapters 4, 11, and 12. These establish that the single most impactful lever for AI success—organizational readiness, consistently the lowest-scoring success factor at 50–60%—is the one most within leadership’s direct control. Cross-functional coordination connects architecture, governance, and business adoption into a single operating model, preventing the fragmentation that limited every prior enterprise technology wave.

Together, these elements define the platform approach this chapter advocates. The difference between organizations that remain trapped in pilot purgatory and organizations that successfully operationalize agentic AI at scale is not the sophistication of their models but the maturity of their architecture, data discipline, governance, and organizational design. The single most impactful lever—organizational readiness, currently scoring 50–60% against technical factors scoring 80–95%—is the one most within leadership’s control. It requires a named owner, typically the CAIO, and a cross-functional coordination mechanism established from day one.

The Path Forward

The GenAI ROI crisis is not a failure of AI, but of the first deployment paradigm. Chatbots and copilots delivered incremental assistance when enterprises expected transformative process automation. Agentic AI resolves this structural mismatch by shifting from passive text generation to autonomous workflow execution. However, agentic AI will not succeed simply because the underlying technology has become more capable. It will only succeed when it is deliberately deployed within the strict architecture, governance, and organizational design frameworks that this book describes.

The organizations that win this transition will not be the ones with the flashiest models. They will be the ones that successfully operationalize governed, scalable action across their complex, end-to-end business workflows. The technology is thoroughly ready. The architecture is clearly defined. The only remaining question is whether leadership has the resolve to invest in the critical foundations of content, orchestration, governance, and organizational change required to transform isolated pilots into enterprise platforms, and expensive experimentation into durable enterprise value.

Executive Implications

CAIO

Scaling AI requires coordinated governance and adoption across the enterprise. The CAIO aligns architecture, policy, and business leadership so the shift from pilots to platforms becomes an operating model transformation.

CIO

AI value does not scale through point solutions. The CIO stewards the enterprise agentic platform, providing the architecture and coordination needed to enable intelligent workflows across the organization.

CFO

The GenAI ROI crisis is structural, not technological—capital should shift from scattered pilots to platform investments that redesign workflows and produce measurable P&L impact.

CHRO

Moving from pilots to platforms demands organization-wide change management, new human–agent operating roles, and leadership-driven adoption—not just training on tools.

CDO

Data quality and enterprise memory are the gating factors for scale; without governed, lineage-rich content foundations, agentic deployments amplify risk instead of value.

COO

Operational gains come from end-to-end workflow redesign and orchestration, not isolated automation; platforms enable durable cycle-time reduction, reliability, and scalability.

In the next chapter, we present a phased adoption roadmap that translates this comprehensive architectural vision into a highly practical 24-month transformation program.

Chapter Twelve

A Phased Adoption Roadmap

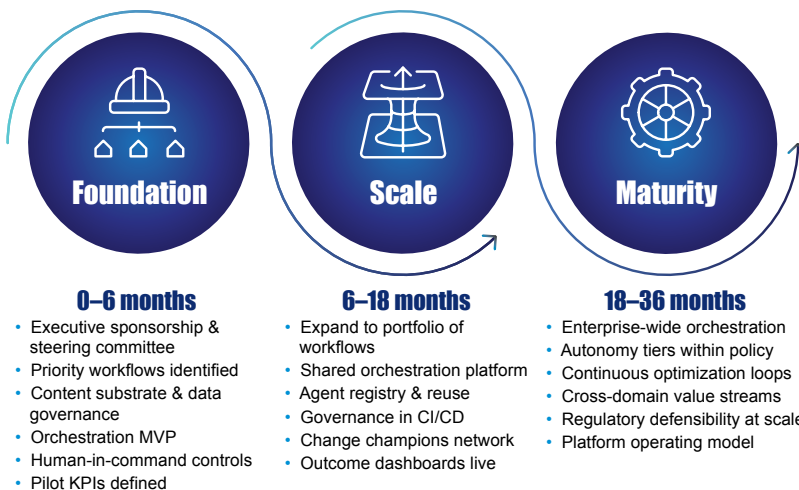
Chapter 11 made the case that value from agentic AI does not emerge from isolated pilots, better models, or ad hoc experimentation—it emerges from deliberate operating model transformation. This chapter translates the platform shift into a sequenced transformation program—foundations, pilot, scale, and optimization.

As we established in Chapter 10 of the first book in this series, there are five key principles for establishing the right conditions to support and manage the deployment of agentic AI applications. These include consideration of the organizational model for deployment, developing the applications, addressing collaboration between human and agentic workforces, and performance management and measurement. In this chapter, we build on this high-level framing to provide a detailed, practical guide for establishing the foundations of agentic AI and then scaling it to maturity over a 24–36 month transformation horizon.

The guidance in this chapter is grounded in real-world deployments—within our own organization and across customer implementations globally—including lessons learned from missteps along the way.

We'll cover that in depth in this chapter, including how to:

1. Build the right program sponsorship and steering committee
2. Address required skills and roles
3. Drive cultural change
4. Choose the appropriate technology partners
5. Identify IT system gaps and bottlenecks
6. Identify end-to-end workflows and business processes for streamlining
7. Select your first project as the foundation
8. Measure and fine-tune outcomes
9. Develop the capability to deliver multiple streams in parallel
10. Follow a prioritized transformation roadmap



1. Build Program Sponsorship

Every agentic AI program begins with sponsorship. Even small initiatives lay the foundation for a cross-enterprise platform that will reshape how work is done. Strong executive sponsorship is required from technology leadership (CIO), people leadership (CHRO) to manage change, and the business owner most impacted (e.g., CFO for finance use cases). Aligning these stakeholders clarifies outcomes, ensures sustained support, and provides the governance needed to resolve issues as they emerge.

It is also helpful to establish a formal steering committee. Their mandate should include:

- Defining enterprise-level objectives
- Setting governance and risk guidelines
- Prioritizing business processes and workflows
- Resolving cross-functional conflicts and roadblocks
- Monitoring business case and value realization
- Approving specific autonomy tiers for agents

Agentic AI programs fail when accountability is not clear. They succeed when ownership is explicit. The steering committee can ensure that project ownership and accountability are well defined and executed according to plan.

2. Address Skills and Roles

Before the agentic AI transformation program starts, an exercise must be undertaken to truly understand the breadth of roles and skills across the organization. Agentic AI will complement existing processes within the organization, in some cases replacing certain tasks.

At OpenText, we began by analyzing the company's role groups and zeroing in on a few that offered strong opportunities to build our agentic AI foundation. In our case, this included task-based and decision-making roles, as we felt that in the first phase, it was important to understand both. We also took the conservative view that for the first 9-12 months, each of these roles would have a human in the loop. This was especially critical for decision-making roles where maturity in technology and consistency in business outcomes are required to trust agentic AI in that capacity.

Start with junior or associate-level roles and clearly distinguish task execution from decision authority. To simplify adoption, define job descriptions for both agents and humans—clarifying scope, responsibilities, and expected outcomes from the outset, as illustrated in the following graphic, which compares both human and agentic AI HR roles.

Human Job Description: HR Operations Specialist

Role Summary

The HR Operations Specialist is responsible for managing and resolving employee HR tickets, with a focus on benefits selection, onboarding, and policy guidance. This role works in tandem with an Agentic AI assistant to ensure timely, accurate, and personalized support.

Key Responsibilities

- Review and validate employee-submitted benefit selection requests.
- Provide personalized guidance based on employee eligibility, location, and role.
- Escalate complex or exception-based cases to HR leadership.
- Collaborate with the AI agent to monitor ticket queues and prioritize urgent cases.
- Audit benefit selections for compliance with internal policies and regulatory requirements.
- Train and calibrate the AI agent by reviewing its recommendations and feedback loops.

Skills & Qualifications

- 3+ years in HR operations or benefits administration.
- Strong understanding of enterprise HRIS systems and benefits platforms.
- Excellent communication and decision-making skills.
- Comfortable working alongside AI agents and digital workflows.

Collaboration with AI Agent

- Oversee and approve benefit recommendations generated by the agent.
- Provide context and nuance for edge cases the agent flags as ambiguous.
- Participate in continuous improvement of agentic workflows and training data.

Agentic AI Agent Job Description: HR Ticket Resolution Agent

Role Summary

The Agentic AI Assistant is designed to autonomously process HR-related tickets, with a primary focus on benefits selection. It operates in close coordination with human HR specialists to ensure accuracy, compliance, and employee satisfaction.

Key Responsibilities

- Automatically classify and route incoming HR tickets using natural language understanding.
- Retrieve and analyze employee data (e.g., tenure, location, job level) to recommend appropriate benefits packages.
- Generate personalized benefits summaries and FAQs for employees.
- Flag tickets requiring human judgment or policy exceptions.
- Learn from human feedback and update decision models accordingly.
- Maintain audit logs and traceability for all actions taken.

Capabilities

- Integrate with enterprise HRIS, payroll, and benefits systems using secure, auditable and approved APIs.
- Use policy documents and historical ticket data to inform decisions.
- Operate 24/7 with real-time response capabilities.
- Continuously improve via feedback-driven calibration and model tuning.

Collaboration with Human

- Send benefit recommendations to the HR Operations Specialist for approval.
- Receive feedback on rejected or modified recommendations to refine future outputs.
- Alert the human to anomalies, missing data, or policy conflicts.

Governance & Oversight

- All actions are logged, and subject to review by automated Audit Agents and HR specialist teams.
- The system operates with autonomy but under continuous monitoring and auditability and undergoes periodic audits to validate compliance and model performance.

It is also worth noting that agentic AI transformation programs will introduce the need for new roles within the organization to support them. These include:

- **Chief AI Officer (CAIO)** as the primary architect of AI governance, working in concert with the CIO and CISO to define the guardrails that individual owners must enforce
- **Human Oversight Owner** for each agent, operating within a framework architected by the Chief AI Officer (CAIO)
- **AI Governance Manager or Supervisor** to oversee policy development and enforcement, compliance, and auditability
- **Agent Architects** responsible for designing the orchestration approach and boundaries for agents, as well as workflow logic
- **Prompt and Interaction Engineer** who is responsible for defining how agents will interpret the context and human input in a repeatable and reliable approach
- **AI Ops Engineer** to conduct real-time monitoring of agent performance and detect any anomalies in behaviors
- **Product Owners** within each business unit, responsible for the specific business KPIs and outcomes

The transformation will require traditional roles to evolve. Business analysts will shift from documenting processes to re-engineering workflows, while data engineers move toward real-time integration and orchestration. Supervisory and management roles will also change—from overseeing individual tasks to governing outcomes, performance, and risk across intelligent workflows. The role of the Chief AI Officer (CAIO), for example, is emerging as the owner of AI outcomes—aligning investments to business strategy, orchestrating cross-functional execution, and ensuring governance, risk management, and responsible innovation keep pace as AI systems move from advisory support to autonomous action.

3. Drive Cultural Change

As with any transformation, the importance of change management should not be underestimated. Introducing agentic AI is one of the most profound shifts an organization will undertake—not because it adds new tools, but because it changes how work is executed, supervised, and governed. Technology value does not scale on model capability alone; it scales when workflows, roles, and operating mechanisms are redesigned to support a hybrid human-AI workforce.

Treat adoption as an operating-model program, not a communications plan. The practical implication is that change management must be wired directly into delivery: pilot selection, workflow redesign, governance design, enablement, and measurement. This is especially true for agentic systems, where execution moves from “suggest” to “do,” and trust becomes inseparable from oversight, auditability, and clear decision rights.

Build momentum with champions and “workflow-native” wins. Identifying change champions—people who are already embracing automation and are willing to help shape new ways of working—is critical. Even when a use case is not the largest immediate business opportunity, it can be worth prioritizing if it accelerates adoption, builds fluency, and normalizes human-in-command patterns. At OpenText, we began with HR-relevant use cases because the HR organization combined strong technology literacy with enterprise-wide influence—making them natural partners for cross-company change leadership.

Leaders should communicate a consistent narrative that aligns both human confidence and governance discipline:

- 1. Agents are there to augment, not eliminate, human expertise.** We introduce autonomy gradually using human-in-command gates, especially in decision-heavy processes. This also reduces anxiety by making human oversight explicit.
- 2. Governance and transparency are scaling enablers, not brakes.** If you want autonomy at enterprise scale, you need auditable execution, clear escalation paths, and policy enforcement built in from the start. As explored in Chapter 11, analysts increasingly frame this “control plane” capability as the difference between pilots and platforms.
- 3. Early wins will be measurable and visible.** Adoption increases when teams see proof in cycle time, service levels, quality, and reduced rework—not just enthusiasm and anecdotes. Tie recognition and advancement to these outcomes and invest in structured upskilling so new opportunities are real, not rhetorical.

In practice, cultural resistance declines when employees experience agents removing repetitive work while preserving (and often elevating) human judgment. The goal is not to make people “use AI.” The goal is to redesign work so that human effort becomes the exception pathway—reserved for ambiguity, risk trade-offs, and accountability—while agents handle predictable execution under policy.

Finally, change management must include concrete readiness mechanisms: role-based training, “how we work with agents” playbooks, escalation norms, and clear ownership of outcomes. This matters because agentic adoption is accelerating rapidly across enterprise software, which compresses the window organizations have to learn safely and institutionalize new norms.

4. Choose Technology Partners

The market is full of new and exciting AI technologies, and it is easy for employees and teams to get excited about “technology A” or “technology B.” The pace of innovation in the AI space will continue, so understanding the expected outcomes is more important than adopting the technology of the day.

Selecting technology that does not lock your organization in to a proprietary solution and offers the broadest flexibility in the future is critical. Equally important is ensuring that the technology meets your organization’s privacy and security standards, and its requirements for data sovereignty. There is no perfect technology, but selecting vendors with technologies that complement your organization’s strengths and weaknesses is key to driving adoption. Technology should be an enabler for agentic AI—it is not the sole solution to delivering on the promise.

According to Forrester:

53% of enterprise AI decision-makers claim their tech organization is primarily responsible for leading the AI efforts in their organization...and just 7% [point] to business functions.⁵⁹

5. Identify IT System Gaps and Bottlenecks

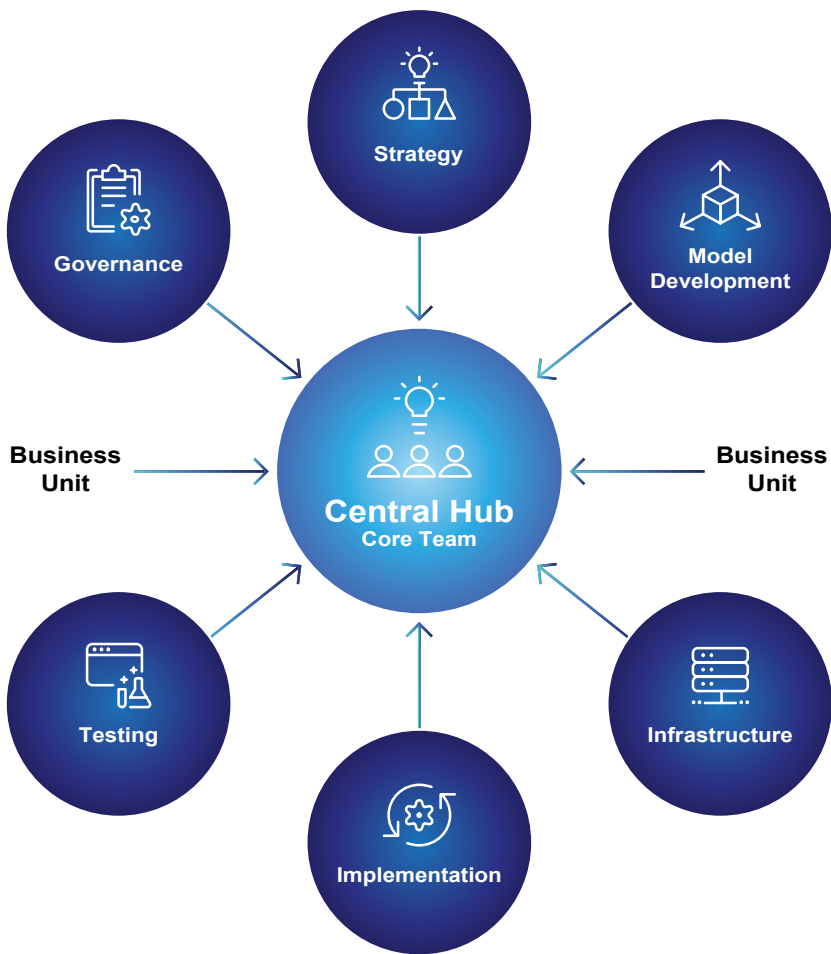
In most organizations today, the CIO and CTO remain the de facto owners of AI strategy, execution, and governance. This is consistent with industry research showing that a majority of enterprises still anchor AI leadership within the technology function, with comparatively limited ownership from business leaders. While this model provides early momentum and technical coherence, it also exposes a structural risk as AI initiatives mature: if AI remains perceived as “an IT program,” business adoption, accountability, and operating model transformation tend to stall. Without distributed ownership, AI risks becoming a bottleneck rather than a catalyst—concentrated in a small expert team while demand for capabilities outpaces delivery capacity.

Agentic AI amplifies this tension. As intelligence moves from the edge of the organization into the core of enterprise workflows, the role of IT necessarily evolves. Technology leaders shift from being the sole builders and operators of AI systems to becoming stewards of the enterprise’s agentic platform—establishing standards, reference architectures, security controls, and governance frameworks that enable business units to safely design, deploy, and scale intelligent workflows. In effect, IT becomes the control plane and enablement layer for agentic transformation, rather than the exclusive owner of execution.

This transition requires a deliberate operating model choice. Depending on organizational maturity, enterprises may adopt a centralized Centre of Excellence (CoE), a hub-and-spoke model, or a federated structure for AI delivery and governance. In earlier stages, a centralized CoE provides the discipline required to establish core architectural patterns, toolchains, security practices, and governance mechanisms. Over time, however, this model must evolve. For a more in-depth discussion of operating models, please refer to Chapter 9 in book one of this series.

In our own journey at OpenText, we began with a centralized CoE embedded within IT to build foundational capabilities across strategy, model development, infrastructure, testing, and governance, in close partnership with security. This created early consistency and control at a time when agentic patterns, orchestration frameworks, and governance models were still being defined. However, we recognized quickly that a purely centralized model can unintentionally absolve business units of accountability for outcomes. To counter this, we partnered closely with domain leaders to co-develop joint business cases, align workflows to strategic objectives, and define shared success metrics.

The long-term objective is a hub-and-spoke model: IT retains ownership of the core agentic platform, orchestration layer, security, and governance, while business units take responsibility for workflow redesign, value realization, and operational adoption.



A Hub-And-Spoke Model

Successful AI programs treat AI not as a technology deployment, but as an operating model transformation that requires shared ownership between technology leadership and the business. Identifying IT system gaps and bottlenecks, therefore, is not simply a technical exercise—it is an organizational design decision that determines whether agentic AI becomes a scalable enterprise platform or remains constrained by centralized capacity and fragmented accountability.

6. Identify End-To-End Workflows and Business Processes

At the heart of effective agentic AI deployments is a deep understanding of your organization's business processes. In previous chapters, we discussed how good data and good business processes are the main ingredients for delivering outcomes with agentic AI.

Transforming end-to-end business processes is a multi-year journey and should be approached in deliberate phases. Start by selecting a high-impact but well-bounded workflow—such as procure-to-pay in finance—where value leakage and manual friction are visible. Begin with a rigorous process mapping exercise to establish a shared, end-to-end view of how the workflow actually operates across the organization.

Next, decompose the workflow into roles, tasks, and decision points. This role- and task-level clarity creates the foundation for identifying where agentic capabilities can safely and productively be introduced.

7. Select the Appropriate First Project

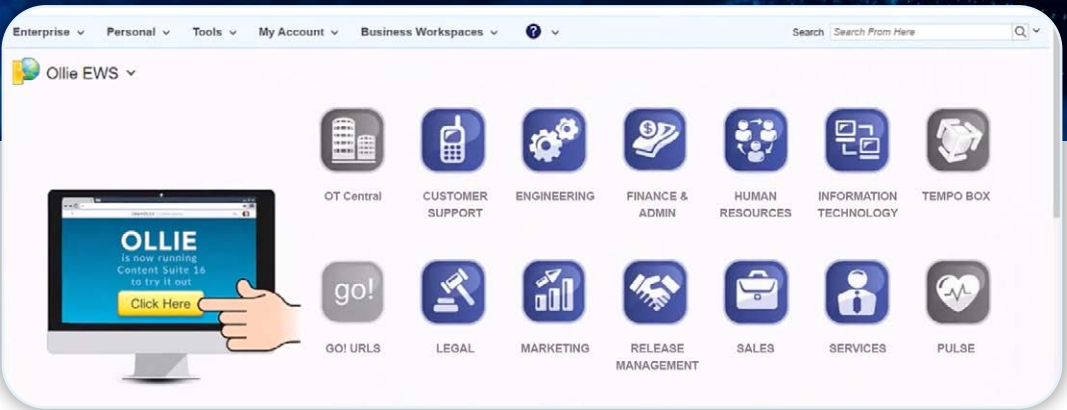
The first project sets the tone for the entire transformation. Early success builds confidence, unlocks sponsorship, and establishes the operating patterns that will carry forward as the program scales. For this reason, the initial use case should be deliberately chosen: a well-defined, end-to-end business process with clear ownership, high visibility, measurable outcomes, and a manageable risk profile.

Executive sponsorship is non-negotiable. While technology leaders often carry formal responsibility for AI strategy and governance, sustained momentum requires explicit business ownership of outcomes. The first project should therefore sit at the intersection of IT enablement and business accountability, with a named executive sponsor who is accountable for realizing value—not just deploying capability. This shared ownership prevents early initiatives from becoming technically successful but operationally irrelevant.

Finally, align expectations from the outset. The purpose of the first deployment is proof of viability, not proof of the organization's ultimate ambition. The goal is to establish the architectural, governance, and operating foundations required for scale: integrating agents into real workflows, validating human-in-command controls, and demonstrating measurable impact. Programs that treat the first project as a controlled production experiment—rather than a showcase pilot—create the conditions for durable, platform-level adoption.

The following case study demonstrates how early, workflow-embedded AI deployments can be deliberately designed to build the organizational, architectural, and governance foundations required for agentic AI at scale. The transition from vision to velocity requires a tangible entry point; for OpenText, this began with their intranet—a deployment that serves as the “phase 1” bridge between traditional productivity and the agentic future by normalizing AI-human collaboration within a governed, everyday interface.

OpenText



As organizations move from experimentation to operationalizing AI at scale, the OpenText intranet (OLLIE) illustrates how early deployments can be deliberately designed to lay the groundwork for agentic AI—without overreaching current maturity. OLLIE embeds AI assistance directly into employees' everyday work environment, providing contextual access to enterprise knowledge at the point of execution. Rather than positioning AI as a standalone interface or isolated tool, OLLIE integrates intelligent capabilities into existing workflows, enabling employees to search, summarize, plan, translate, and generate content without leaving their core systems. This reduces friction in day-to-day work and accelerates decision cycles by keeping knowledge and action tightly coupled.

While OLLIE is not, in itself, a fully autonomous agentic system, it establishes several of the foundational conditions required for agentic AI at scale. By grounding AI interactions in governed enterprise content and embedding them within operational workflows, OLLIE reinforces the principles of contextual intelligence, trusted enterprise memory, and workflow-integrated AI described earlier in this chapter. Employees are not simply consuming AI outputs; they are learning to collaborate with AI as part of how work gets done.

Equally important, deploying AI through a trusted intranet platform creates a controlled environment for change management. It allows the organization to establish patterns for human-in-command oversight, content governance, identity and access control, and user adoption—capabilities that become non-negotiable as autonomy increases. In this way, OLLIE functions as a “training ground” for the organization itself, building literacy, confidence, and operational muscle memory for more advanced agentic workflows. By embedding AI into the fabric of daily work through the intranet, OpenText has created a foundation upon which more autonomous, orchestrated agentic capabilities can be safely introduced over time—shifting AI from a productivity layer to an operating capability that scales.

Executive Decision Framework: Pilot or Platform?

For executives assessing their AI portfolios, the following framework provides a clear diagnostic—distinguishing scalable enterprise platforms from perpetual pilot purgatory.

Dimension	Pilot Mindset	Platform Mindset
Starting point	Which tasks can AI accelerate?	Which workflows should be redesigned?
Architecture	Point solutions per use case	Shared orchestration and content substrate
Governance	Added after deployment	Embedded from ideation
Data strategy	Whatever the model can access	Governed enterprise memory with lineage
Measurement	Demo worked; users like it	Business KPIs: cycle time, quality, cost
Agent management	One-off builds	Registry of approved agents, tools, workflows
Human role	Optional reviewer	Defined accountability, escalation, oversight
Scaling approach	More pilots	Portfolio expansion on shared platform
Organizational change	Training on the tool	Workflow redesign with change management
Outcome	Impressive demos, modest impact	Measurable, defensible operating value

Organizations that find themselves in the “Pilot Mindset” column need not despair. Every enterprise that has successfully scaled agentic AI began with pilots. The difference is discipline: treating pilots as production-bound experiments with clear learning objectives, measurable outcomes, and defined pathways to scale—rather than as impressive demos with no operational future.

8. Measure and Fine-Tune Outcomes

Whether an agentic deployment succeeds or fails is determined in design but measured in production. For this reason, outcome measures must be defined before deployment and wired directly into live workflows. The roadmap requirement is straightforward: define outcome KPIs before deployment, instrument the workflow end to end, and operationalize continuous tuning and retirement criteria from day one. Measurement is not a reporting exercise; it is part of the operating model for agentic AI.

Teams should establish clear baselines for the workflows being augmented—cycle time, error and fallout rates, cost-to-serve, SLA performance, and escalation volumes—so that improvements can be credibly attributed to agentic execution. Instrumentation must span the full workflow, from signal ingestion through orchestration to downstream system actions, enabling real-time observability of performance, exceptions, and policy adherence. Dashboards should be owned by named business and technology leaders, with explicit accountability for outcomes, not just usage.

Equally important is a defined remediation and decommissioning model. Not every agent will perform as intended in production, and some will outlive their usefulness as workflows evolve. High-performing programs treat underperforming agents as learning assets: diagnosing root causes, determining whether remediation is warranted, and retiring agents that no longer meet performance, risk, or governance thresholds. This operationalizes a culture of continuous learning that is essential to scaling safely.

Chapter 13 provides the measurement and defensibility framework in full; here, the requirement is simple: define outcomes up front, instrument them end-to-end, and operationalize remediation.

9. Deliver Multiple Streams in Parallel

Once the foundational capabilities are in place, the transformation can shift from deliberate build-out to controlled acceleration. Scaling agentic AI requires more than adding projects; it requires standardizing how intelligence is designed, deployed, and governed across the enterprise. Reusable agent components, shared orchestration patterns, and common design principles become the multiplier that enables parallel delivery without fragmenting the architecture.

This model reinforces the hub-and-spoke operating structure described earlier. The AI CoE serves as the hub—owning the reference architectures, autonomy tiers, lifecycle standards, security controls, sovereign deployment patterns, and shared services that form the enterprise control plane. The domain-specific agile squads operate as the spokes—embedded within business units and value streams, applying these standardized patterns to real workflows. Rather than centralizing execution, the CoE centralizes discipline. This separation allows multiple streams to advance in parallel while preserving architectural integrity, governance consistency, and interoperability across the expanding estate of agents. Acceleration becomes structured rather than chaotic, because every spoke builds on the same governed foundation established by the hub.

Crucially, value realization must be coordinated centrally. Portfolio-level prioritization, funding decisions, and performance tracking should remain visible to the CoE to ensure that parallel initiatives reinforce a coherent operating model rather than competing for resources or drifting into siloed implementations. This balance—federated execution with centralized standards and governance—is what allows agentic AI programs to scale with speed, consistency, and control.

10. Follow a Prioritized Transformation Roadmap

A successful program requires a prioritized roadmap. This may differ based on your organization, but a simple framework includes:

Step 1 – Building the Foundation for Transformation at Scale (Pre-Pilot)

- Establish program steering committee
- Define the AI Governance framework
- Complete data readiness assessment

- Address any IT capability gaps (e.g., identity and access management)
- Ensure API readiness for agentic deployment

Step 2 – Pilot Deployment

- Select a business process or workflow for pilot deployment
- Establish a clear accountability model
- Design for operations, observability, and audit
- Ensure human-in-command oversight model

Step 3 – Scaling Beyond Pilot

- Establish domain-specific squads and begin domain-by-domain rollout
- Ensure shared services and technology components are ready for broader deployment
- Set clear agent and orchestration standards for common adoption
- Identify cross-domain workflow integration priorities

Step 4 – Measuring and Optimizing

- Continue performance tuning to evaluate outcomes
- Refine models as needed
- Ensure cross-functional agent collaboration
- Complete ongoing measurement and roll out the KPI framework

This roadmap follows the platform principles established in Chapter 11—workflow-first, governed enterprise memory, shared orchestration, built-in trust, and adoption by design.

From Pilot to Agentic Enterprise

Your agentic AI transformation program will be successful if:

- ✓ Governance and innovation move at pace, together
- ✓ Humans remain in the loop
- ✓ Infrastructure supports real-time orchestration
- ✓ Agents operate within defined autonomy tiers
- ✓ Value is measurable and defensible
- ✓ Scaling is handled systematically

Your Agentic AI Genome will be built through disciplined orchestration of people, processes, platforms, and policy. The agentic AI transformation is about redesigning the enterprise to think, act, and adapt with intelligent systems—responsibly, securely, and at scale.

Executive Implications

CAIO

Enterprise-scale adoption requires coordinated governance across technology, compliance, and business leadership. The CAIO ensures each phase of adoption delivers governed, auditable outcomes—not just technical capability.

CIO

Agentic AI scales only when it is treated as a platform transformation, with shared orchestration, governance, and delivery standards enabling parallel execution across the enterprise.

CFO

Value is realized not from isolated pilots, but from a sequenced investment roadmap that ties agentic deployment to measurable business outcomes, cost-to-serve reduction, and defensible ROI.

CHRO

Sustainable adoption requires redesigning roles, skills, and performance models so that human–agent collaboration becomes the operating norm, not an exception.

CDO

Enterprise data readiness and governed content substrates are prerequisites for scale; without trusted, accessible enterprise memory, agentic systems cannot move safely from pilot to platform.

COO

Operational impact emerges when end-to-end workflows are deliberately redesigned for agentic execution, shifting the enterprise from task automation to orchestrated, outcome-driven operations.

In the final chapter in Part 4, we'll delve into measuring the value and defensibility of agentic AI adoption.

Chapter Thirteen

Measuring Value and Defensibility

Chapter 13 closes the loop on the agentic AI transformation by addressing the question that ultimately matters most to leadership: how to prove that agentic AI is delivering durable business value—safely, defensibly, and at scale. It outlines how to define and measure ROI, quantify risk reduction, and establish the confidence that enables the deployment of agentic AI.

Agentic AI marks an important shift in enterprise value creation. Unlike first-generation AI systems that search, summarize, and inform decision-making, agentic AI systems execute decisions. They orchestrate workflows, reconcile data, enforce policy, and initiate actions across complex environments. As discussed in previous chapters, this transformation requires governance, sovereignty, and human-in-command design. But governance alone is not enough.

Leaders must be able to answer the fundamental question: Is it working, and are we generating real enterprise value?

Today's AI deployments are measured not only in financial returns but also in risk reduction and operational resilience. Enterprises must move beyond narrow productivity metrics and toward a multi-dimensional model of agentic performance measurement.

We provide a practical framework for operationalizing trust: embedding observability, accountability, and policy enforcement into agent design; wiring performance measurement directly into orchestration; and institutionalizing governance as a continuous control plane rather than a periodic review. The result is a model for scaling agentic AI with confidence—where value is measurable, risk is managed by design, and autonomy expands in lockstep with accountability.

Rethinking ROI for Agentic AI

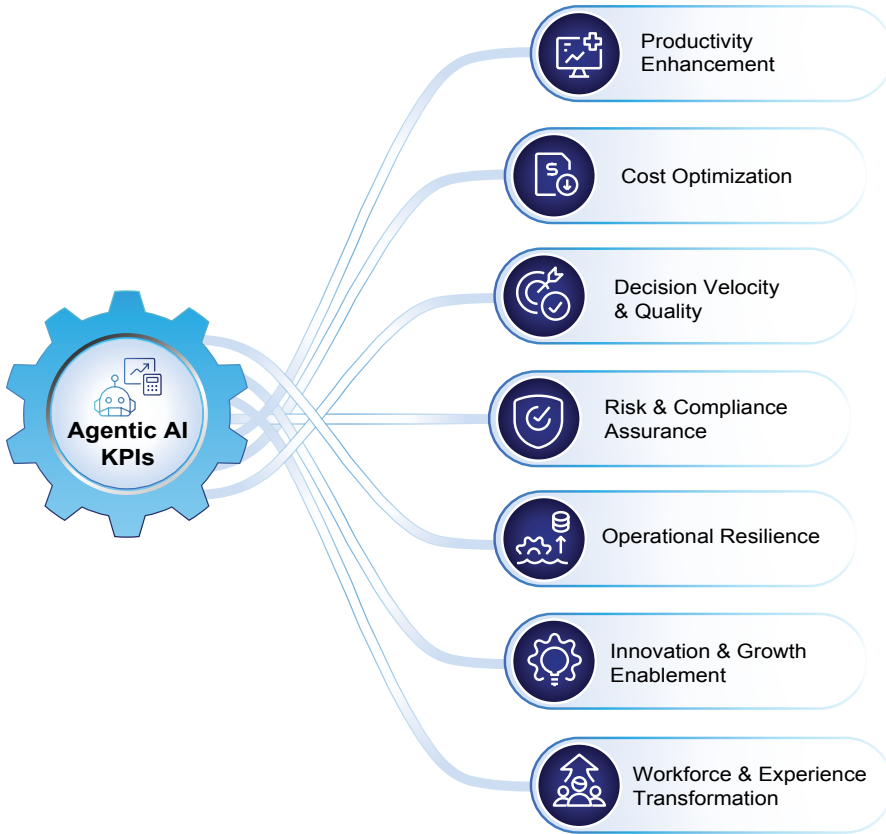
Traditional technology business cases are built around replacing a specific system or technology, automating tasks, or improving specific KPIs. Agentic AI forces a rethink of the business case because it impacts end-to-end business processes, touches multiple systems, introduces new decision-makers into the process, and alters the overall velocity of operations.

For agentic AI, ROI is comprised of the following components:

- Headcount savings from advanced automation
- Faster time-to-revenue (optimized processes, onboarding)
- Reduction in revenue leakage (improved processes)
- Optimization of network and infrastructure usage

Before agentic AI deployment, it is important to establish a strong baseline of costs to ensure measurements are meaningful. This includes cost per transaction (systems, labor, infrastructure, and operations), error and fallout rates (the share of automated processes that “fall out” of normal automated workflows), time-to-revenue, and penalties and impact of missed SLAs.

There are many other relevant KPIs to consider, and it's crucial to identify which metrics are significant for your organization. Below, we analyze KPIs across seven categories to inform business case development:



Driving Productivity Enhancement and Cost Optimization

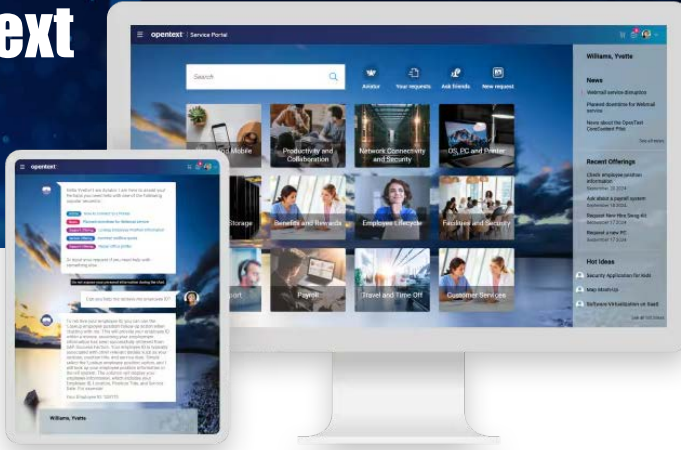
Agentic AI can enhance productivity, drive operational efficiency, and reduce costs. The move from simple automation to end-to-end workflow orchestration enables this. Traditional automation and RPA approaches delivered positive business case outcomes, but agentic AI can serve as a force multiplier on top of them.

Agentic AI enables the orchestration of flows across different systems, interpreting context and executing actions and decisions to drive outcomes. As agentic AI handles repetitive tasks, human teams can shift their focus towards higher-value work, which results in measurable improvements in throughput, lower processing costs, and reduced fallout.

If the ultimate proof of an agentic transformation lies in its ability to move the needle on the P&L, then OpenText provides the definitive blueprint—demonstrating how a modernization-first approach to service management can turn a \$1B productivity goal into a measurable reality by reducing ticket volumes by 70%.

Durable value is unlocked not by deploying more powerful models in isolation, but by first modernizing processes and data, then layering autonomous capabilities into governed production workflows.

OpenText



As enterprises begin operationalizing agentic AI, service management is emerging as one of the most natural entry points for embedding intelligence into real workflows. At OpenText, service management has been modernized as part of a broader move toward AI-enabled operations—integrating automation, intelligent routing, knowledge retrieval, and workflow orchestration directly into the way employees request, receive, and resolve support. Rather than treating AI as a standalone assistant layer, service management embeds intelligence into the execution fabric of IT and enterprise services, enabling faster triage, more consistent resolution paths, and tighter integration between systems of record. This shift reflects a broader pattern in agentic adoption: early value is created where AI can sense operational context, reason over enterprise data, and participate directly in end-to-end workflows under governance and policy controls.

Service management at OpenText was the first implementation of agentic AI within the enterprise and formed a core pillar of the company's \$1B savings program to drive productivity and efficiency at scale. The transformation began with service management modernization—investing first in process discipline and data quality to establish a reliable operational foundation. This initial phase delivered a 30% reduction in first-line support requests by improving workflow design, knowledge accessibility, and automated handling of routine issues. With this foundation in place, OpenText expanded toward more comprehensive agentic capabilities across service workflows, increasing the overall reduction in ticket volumes by 70%.

We now have a single helpdesk portal for employees, with embedded AI that knows the corporate data inside and out. No endless emailing or bouncing between platforms, and everything in one place. The result has been a fundamental shift in how employees access and receive support—moving from reactive, manual service delivery to orchestrated, AI-enabled service management—now supporting more than 500 teams across the organization.

This progression illustrates a core principle of the agentic adoption roadmap: durable value is unlocked not by deploying more powerful models in isolation, but by first modernizing processes and data, then layering autonomous capabilities into governed production workflows. When performance measurement is wired directly into orchestration, autonomy expands in lockstep with accountability, turning a cost center into a high-velocity engine of enterprise value.

The Value Realization Path



Enhanced ROI through Decision Velocity and Quality

One of the key ROI advantages coming from agentic AI is increased decision velocity—the speed at which accurate, policy-driven decisions move through an organization. As decisions shift from sequential, human-bound handoffs to orchestrated, machine-speed execution, enterprises realize tangible gains: faster approvals, shorter cycle times, fewer bottlenecks, lower error rates, and improved SLA performance.

While this drives material benefits, it can be difficult to measure. To quantify its impact, organizations must first establish a baseline view of where decisions stall today—across workflow bottlenecks, approval queues, SLA breaches, and downstream error or rework rates. Without this operational baseline, improvements in decision velocity remain anecdotal rather than defensible, making it harder to translate agentic performance into credible business value.

From a model perspective, think of it in the context of the table with hypothetical stats below:

Metrics for a specific business process	Before	With Agentic AI	Improvement
Average Decision Time	5 days	4 hours	97% improvement
Human Handoffs	7	3	58% reduction
Fallout Rate	15%	5%	67% reduction
Error Rate	3%	1%	67% reduction

Measuring ROI of Risk and Compliance Assurance

When architected and governed correctly, agentic AI becomes a mechanism for reducing operational, legal, and regulatory risk—embedding policy enforcement, auditability, and control directly into execution. When implemented poorly, however, it can amplify risk by accelerating flawed decisions, expanding compliance exposure, and obscuring accountability. Governance, therefore, is not optional. It is the condition under which agentic AI becomes a risk mitigator rather than a risk multiplier.

From an operational risk perspective, agentic AI reduces variance in execution by embedding policy enforcement directly into workflows. Autonomy tiers force explicit decision boundaries, shifting compliance from reactive oversight to proactive, structural control. In practice, this translates into lower error rates, fewer exceptions requiring human intervention, reduced policy violations, and cleaner audit outcomes—driving margin improvement and lower operating costs.

From a regulatory and legal risk standpoint, well-governed agentic systems provide end-to-end auditability. Decisions are logged, data lineage is traceable, and policy and model versions are controlled by design. This materially reduces compliance exposure by accelerating audit response, lowering remediation effort, and minimizing regulatory findings and escalation events.

While compressed cycle times provide a compelling headline for ROI, in the high-stakes world of global finance, speed is a liability unless it is anchored by absolute defensibility; this leading institution in the following case study illustrates how decision velocity is transformed from a metric of efficiency into a pillar of institutional trust by embedding governance directly into the orchestration of billions of real-time transactions.

As financial institutions operationalize AI, few environments demand more measurable defensibility than global banking. At a leading financial institution, AI adoption evolves from isolated modeling use cases to enterprise-wide workflow integration—particularly across fraud detection, compliance monitoring, and contract intelligence.



A Global Financial Institution



ROI without defensibility is not durable. In highly regulated environments, productivity gains are measurable only if error rates, bias exposure, and policy compliance are equally measurable.

Unlike advisory AI tools, the bank's AI systems for governance participate directly in high-stakes workflows: flagging anomalous transactions, triaging suspicious activity reports, and accelerating loan underwriting decisions. The value equation is measurable across three dimensions:

- **Decision Velocity:** In fraud detection, investigation times dropped from days to minutes by shifting from manual review queues to orchestrated machine triage.
- **Risk Reduction:** AI systems embedded in regulatory reporting have improved audit traceability, providing regulators with defensible proof of how decisions were made.
- **Operational Resilience:** By analyzing billions of transactions in near real-time, the bank has materially lowered fraud losses while reducing the fallout rate of manual exceptions.

Establishing the Measurement Baseline

To quantify the shift from manual to agentic execution, the institution could use the following performance framework:

Decision Metric	Manual/Legacy Baseline	Agentic Orchestration	Strategic Value
Investigation Time	3-5 Days	15-30 Minutes	Increased Liquidity/ Customer Trust
Audit Traceability	Periodic/Sampling	Continuous/100%	Reduced Regulatory Fines
Human Handoffs	High (Multi-department)	Low (Exception-only)	Lower Operating Cost/Error Rate

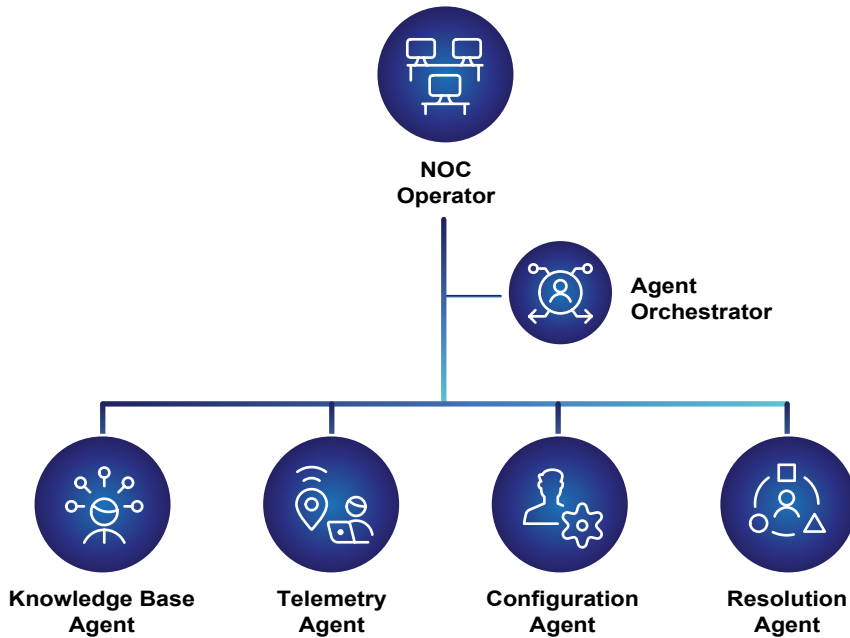
This example underscores a critical principle: the value equation for agentic AI is a combination of Business Value + Risk Reduction + Operational Resilience. By moving governance from a periodic review to a continuous control plane, the bank can ensure that its 97% improvement in decision speed does not come at the cost of its 100% requirement for legal certainty. Their architecture serves as a blueprint for the Agentic Genome, where every autonomous action carries its own decision-making proof, ensuring the system remains as auditable as it is fast.

ROI through Operational Resilience

The impact of operational resilience is broad reaching and underestimated in ROI measurement. To illustrate, recall the Facebook outage in 2021 which also impacted Instagram and WhatsApp. Billions of people were locked out for 6 hours.⁶⁰ Beyond the inconvenience to its global users, the outage cost Facebook in lost revenue and brand damage. This is why operational resilience is critical, and establishing the true ROI for this challenge can be difficult.

Operational resilience measures are typically measured in terms of mean time to detect (MTTD) and restore (MTTR) as well as system failover response times. SLAs are set with recovery time objectives (RTO) and service availability metrics.

Agentic AI improves these metrics through real-time monitoring and dynamic workflow rerouting. Agentic systems can maintain operational continuity during infrastructure instability without requiring human intervention. Even in cases where agentic AI is not driving operational decision making, there is a strong business case for amplifying the reach of operations engineers through agentic AI. In early implementations, we see a 5:1 ratio of agents to humans. The diagram below shows how our Network Operations Center (NOC) Operators at OpenText manage a set of operational agents and orchestrators.



Managing Operational Agents and Orchestrators

The Defensibility of Agentic AI

Defensibility of AI is the ability to explain, audit, and legally stand behind the actions of an agentic AI system. While traditional AI systems were more advisory, with humans taking the advice and deciding whether to act, agentic AI shifts that responsibility onto the agent.

As a CIO deploying AI technology, it's important to be able to answer the following questions:

Why did the agent act?

What data did it use?

Which policy governed the action?

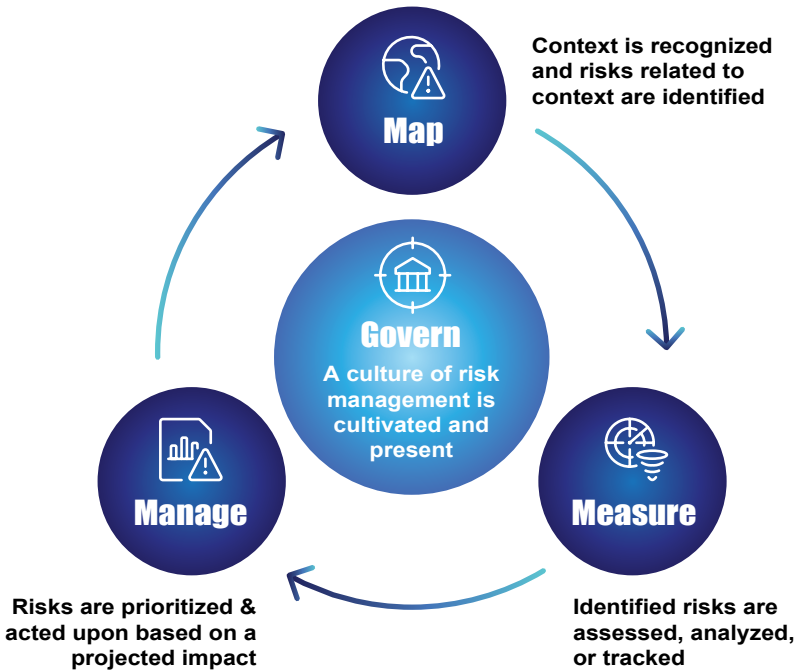
Was the data accurate and lawful to use?

Could the decision be reversed?

Does it comply with applicable regulations?

There is not one single defensibility framework in the industry today, but there are several risk management frameworks that can provide guidance and assistance in defining your defensibility framework.

The National Institute of Standards and Technology (NIST) AI Risk Management Framework (AI RMF 1.0) is a voluntary, flexible, sector-agnostic guidance framework designed to help organizations systematically govern, map, measure, and manage the risks associated with AI systems throughout their lifecycle. Rather than prescribing rigid controls, the framework defines four core functions—Govern, Map, Measure, and Manage—that together provide a structured model for identifying risk factors, establishing accountability, assessing performance and impacts, and implementing risk mitigation and oversight in context-specific ways. By embedding risk management activities into every stage of AI design, deployment, and operation, the AI RMF promotes trustworthy, transparent, and accountable AI systems that align with organizational values, regulatory expectations, and societal norms.⁶¹



NIST AI Risk Management Framework⁶²

The NIST framework supports defensibility of AI through its focus on documentation, risk assessment, model evaluation, and ongoing lifecycle governance.

Another framework is the **EU Artificial Intelligence Act**, which is the European Union’s landmark regulatory framework governing the development, deployment, and use of AI systems across all sectors within the EU market, designed to balance innovation with safety, fundamental rights, and trust. It adopts a risk-based approach that categorizes AI into levels—from prohibited practices and high-risk systems to minimal or no-risk applications—and applies corresponding obligations for transparency, conformity assessment, and governance, including specific provisions for general-purpose and systemic-risk models.

The Act also establishes enforcement mechanisms, conformity requirements, and governance structures, including national supervisory authorities and an EU AI Office. The Act's phased implementation timelines span 2026–2027, while reinforcing human oversight, accountability, and legal certainty in AI adoption.⁶³

As your organization implements its own defensibility framework, it is important to keep in mind the following levels of proof required for advanced agentic AI deployments:

1. **Data Quality Proof** – ensures that your organization has traceability to the system of record for the data and that the data used in decision-making is understood.
2. **Decision-Making Proof** – ensures traceability of the inputs received during the decision-making process and the outputs generated. In this case, the policies that are applied should also be traceable.
3. **Autonomous Action Proof** – a requirement for more advanced autonomous workflows, where it is critical to understand the level of autonomy granted to agentic AI and how that translates into the decision-making process. It is also important to define escalations and how those are triggered.
4. **Legal and Regulatory Compliance Proof** – depending on specific requirements, you may need to establish proof for data residency, regulatory/jurisdictional controls, and more.

Building Blocks for a Measurement Architecture

With all of these considerations in mind, developing a robust measurement architecture requires several building blocks. This includes observability platforms and tools, centralized audit logging, performance dashboards for workflows and business processes, and telemetry for security and performance.

Organizations that have previously invested in KPI models will benefit from a strong baseline. But it is still possible to implement new models and measure benefits. At a minimum, you should plan for real-time dashboards with key metrics, cross-domain performance analytics, and executive KPIs and reporting.

Defining the Value Equation for your organization can help drive culture change and adoption. It can be as simple as a combination of Business Value + Risk Reduction + Operational Resilience, or it can include other factors relevant in your environment. The key point here is that value with increased risk will not drive success. Similarly, if governance increases and value reduces, the transformation will not be a success.

As we have discussed throughout this book, in the Agentic AI Genome, value and governance are inseparable, and ROI without defensibility is not sustainable. So don't just drive deployment of agents. Measure them. Refine them. Defend them. That is the only way to extract their full value for your organization.

Executive Implications

CAIO

Defensibility in agentic AI requires measurable governance. The CAIO establishes the oversight metrics and accountability frameworks that allow leadership to verify agents operate within policy, risk thresholds, and regulatory expectations.

CIO

Defensibility in agentic AI is an architectural responsibility—governance, auditability, and control must be embedded into platforms and orchestration layers, not bolted on after deployment.

CFO

The financial case for agentic AI hinges on defensible outcomes—reduced operational risk, faster decision velocity, and auditable ROI that stands up to regulatory and investor scrutiny.

CHRO

As autonomy increases, human accountability becomes more critical—roles, escalation paths, and trust in human–AI collaboration must be deliberately designed and reinforced.

CDO

Defensible AI depends on defensible data—lineage, quality, and access controls determine whether agentic decisions are explainable, auditable, and regulator-ready.

COO

Operational resilience improves when policy enforcement and auditability are embedded into execution—agentic AI must reduce variance, not introduce new operational risk.

In the next chapter, we explore the possibilities at the frontier of agentic AI at work, where human expertise and machine intelligence converge to redefine roles, accountability, and performance—and unlock new sources of innovation.

Part 5

What Comes Next?

05

Chapter Fourteen

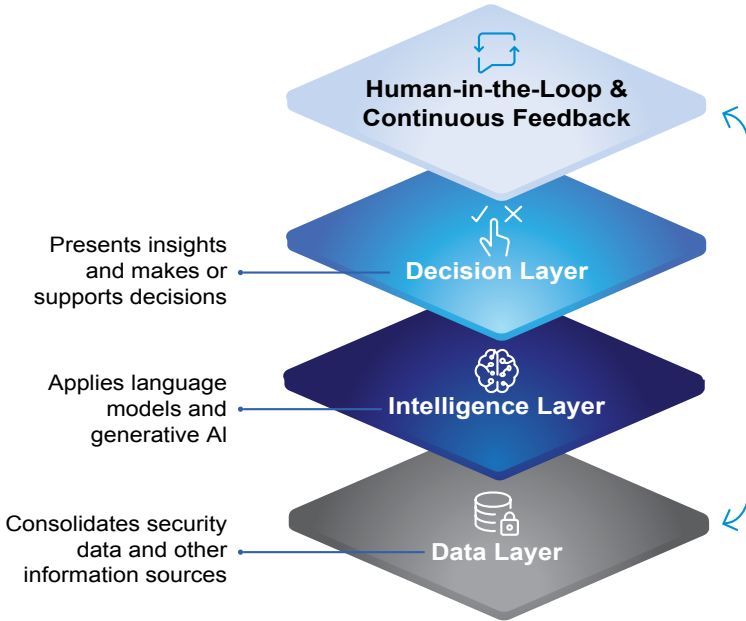
The Future of Human and AI Work

The future of work with the Agentic AI Genome is not a question of humans versus machines. It is a deliberate redesign of enterprise operating models, where agentic AI powers more efficient execution while human judgment remains central. The drive for velocity cannot come at the expense of trust, which is where governance plays a foundational role. In this chapter, we'll explore how organizations can apply the lessons from the Agentic AI Genome to build their own human–AI partnership model. We'll cover how to execute fast without being reckless and how to automate without dehumanizing work.

The Evolution of AI in Work

The nature of work has evolved rapidly from automation to advanced analytics to machine learning and GenAI, and now, to more sophisticated agentic AI. The shift from AI *supporting* decisions to AI *making* decisions means the workforce of the future will be comprised of human and digital employees working in concert with one another. Future org charts will feature both human and digital workers. Humans will set the direction, and agents will deliver the acceleration.

Foundational Layers for Operational AI



Human-in-the-Loop and Agent Collaboration

The agent lifecycle must be managed similarly to the human resource lifecycle. This includes looking at onboarding holistically, clearly defining roles and objectives, understanding governance and access, monitoring and auditing performance, and offboarding workers. Establishing clear job descriptions for both the human and agent can support more effective change management around agentic AI deployment. (See Chapter 12 for greater detail.)

Early enterprise AI delivered predictions, classifications, and generated content. A human asks; the system answers. A human decides; the system supports. This paradigm provided control but it also imposed a ceiling on scale, agility, and real-time responsiveness. At its heart, it is reactive.

Agentic AI breaks this paradigm. It introduces outcome-oriented systems comprised of specialized agents coordinated through orchestration layers. These agents complete end-to-end workflows by retrieving data, applying policies, updating systems, notifying stakeholders, and escalating when required. In this shift, AI moves from tool to workforce. Instances of human-in-the-loop are being replaced by human-in-command—a model where a human defines the boundaries and agentic AI acts autonomously within those boundaries. This is the foundation of the agentic enterprise.

Faster Without Being Reckless

The Agentic Genome is a transformative approach for enterprises, built on a design principle that allows organizations to accelerate their decision-making without veering into recklessness. Traditional workflows often face obstacles and inefficiencies. Agentic AI overcomes this by integrating autonomous, task-focused agents into processes.

However, increased speed can bring increased risk. The Agentic Genome addresses this risk by connecting speed with accountability, ensuring that rapid decision-making is grounded in trust. Instead of treating governance as an afterthought, it's built into the system from the start. A human-in-command approach maintains leadership and accountability while allowing agents to function within specified limits. Autonomy tiers classify decisions by impact, enabling low-risk, routine tasks to be handled independently while humans manage more significant decisions. Policies are embedded to enforce compliance, security, and ethical conduct of agents.

Real-time monitoring and continuous auditability transform governance from a periodic checkpoint into an embedded control plane. Decisions, policy applications, data inputs, and system actions are logged by design, creating end-to-end traceability across workflows. In this model, safeguards do not slow execution; they enable it.

Established governance and accountability frameworks provide structured guidance for integrating audit integrity, transparency, and risk management into AI systems from inception through lifecycle operations. Data privacy, cybersecurity requirements, and other industry-specific regulations define the boundaries for agent operations. When auditability is architected into orchestration and agent design, organizations close the accountability gap while preserving the velocity advantages of agentic execution.⁶⁴

Embedding audit capabilities into your agentic deployment can help to turn decision-making speed into a competitive advantage. By minimizing handoffs and reducing workflow delays, organizations streamline coordination. Throughput increases as agents manage tasks across multiple systems without human intervention. Strategic choices are based on continuously updated data instead of outdated extracts. The outcome is not only quicker operations but also enhanced organizational momentum.

Within the Agentic Genome, speed and trust work in harmony, evolving together as core capabilities of the organization.

To see how embedded accountability actually accelerates delivery, we can look at a National Land and Property Agency that transformed its month-long onboarding bottlenecks into near-instant access through automated governance.



A National Land & Property Agency



A national land and property agency responsible for mapping, boundary management, and maintaining a property registry supports more than 2,000 employees and approximately 20,000 external users—including financial institutions, agricultural stakeholders, and citizens—who rely on secure access to land and property information. As the agency began modernizing its operations and consolidating business units to provide more integrated services, it also explored how intelligent automation and AI systems could improve service delivery and internal efficiency.

The agency was hampered by over 200 applications operating across different platforms, each with its own identity management approach. Without centralized visibility into user access, onboarding new employees required manual coordination across systems and took over a month. Multiple credentials created operational friction while also introducing security and compliance risks. To address these challenges, the organization implemented a unified identity and access management architecture designed to support both human users and automated workflows. Clear governance boundaries were established for automated processes, ensuring that every action taken across systems adhered to policy-defined permissions and audit controls.

With this governance foundation in place, identity services functioned as a control layer, enabling automated processes to retrieve information, execute tasks, and interact with systems of record while maintaining full traceability. Every access event, system interaction, and policy enforcement action could be logged and monitored, creating end-to-end visibility across both human and machine-driven activities. The results were immediate. New employees now receive access to core systems on their first day, with additional systems provisioned within five days. Integration between development and operations teams also allows 98.5% of identity and access issues to be resolved within two days. With consolidated identity governance across more than 200 applications, the organization strengthened information security while significantly improving operational efficiency.

The transformation demonstrates an important principle of the agentic enterprise: autonomous execution can only scale with governance-by-design. By architecting auditability into the very fabric of their identity services, the agency closed the accountability gap while simultaneously unlocking unprecedented velocity. In this model, the safeguards didn't slow the mission; they provided the structural integrity required to move at the speed of an agentic enterprise.

Automated Without Being Dehumanized

The Agentic Genome enables enterprises to automate without becoming dehumanized. This is accomplished by redefining the human role in an AI-first organization. As intelligent agents assume responsibility for high-volume, repetitive, and rules-based workflows, humans shift from task execution to oversight and judgement. Instead of processing information manually, employees define strategy, architect the framework, and approve processes where needed. Rather than managing repetitive workflows step by step, the humans-in-the-loop govern exceptions, validate edge cases, and ensure that automated systems remain aligned with strategy. (For an example, see the Australian Port Authority case study in Chapter 7). In this transition, humans move up the value chain into a leadership position.

In the context of agentic AI, there is a clear distinction between automation and augmentation: “Automation implies that machines take over a human task, augmentation means that humans collaborate closely with machines to perform a task.”⁶⁵ Critically, the Agentic Genome embeds transparency into the underlying AI architecture where human augmentation is paramount. Employees must have clear visibility into how agents reach conclusions, what data is used, and which policies shape the decisions. Explainability and defensibility are foundational requirements. Escalation rights are equally important. When employees can challenge, override, or review automated decisions, autonomy remains balanced by human authority. This ensures that automation augments, rather than replaces, professional judgment.

Protecting human creativity, empathy, negotiation, and ethical reasoning are differentiating capabilities in a competitive enterprise. Automation should remove friction. By eliminating manual work, the Agentic Genome frees up human capacity for innovation, collaboration, and leadership:

“Employees and managers whose basic skills are made redundant by automation could be given the opportunity to gradually build higher-level augmentation skills that remain in demand. This skill-enhancement cycle could also help ‘rehumanize work’ by gradually shifting the focus from repetitive and monotonous tasks to more creative and fulfilling tasks.”⁶⁶

When technology handles precision and scale, and humans can focus on strategy and direction, the enterprise is faster and more efficient, yet fundamentally more human. This is demonstrated in the Roche Diagnostics case study (Chapter 7), where automation and AI covers core testing so the development team can focus on innovation—supporting higher levels of job satisfaction.

Skills for the Agentic Workforce

Earlier in this book we described a fundamental shift in how organizations scale work. In the agentic enterprise, the human workforce becomes a relatively slow-moving stock variable, while the population of digital agents becomes a fast-scaling flow variable. Hiring cycles operate on the scale of months or years, but agents can be deployed, modified, and orchestrated in hours or days. This difference fundamentally changes how organizations compete. As digital execution expands, human roles increasingly concentrate around decision-making, supervision, ambiguity resolution, and accountability, while agents perform the execution, monitoring, analysis, simulation, and coordination that once required large operational teams. The skills required of the workforce therefore evolve alongside the architecture of the enterprise itself.

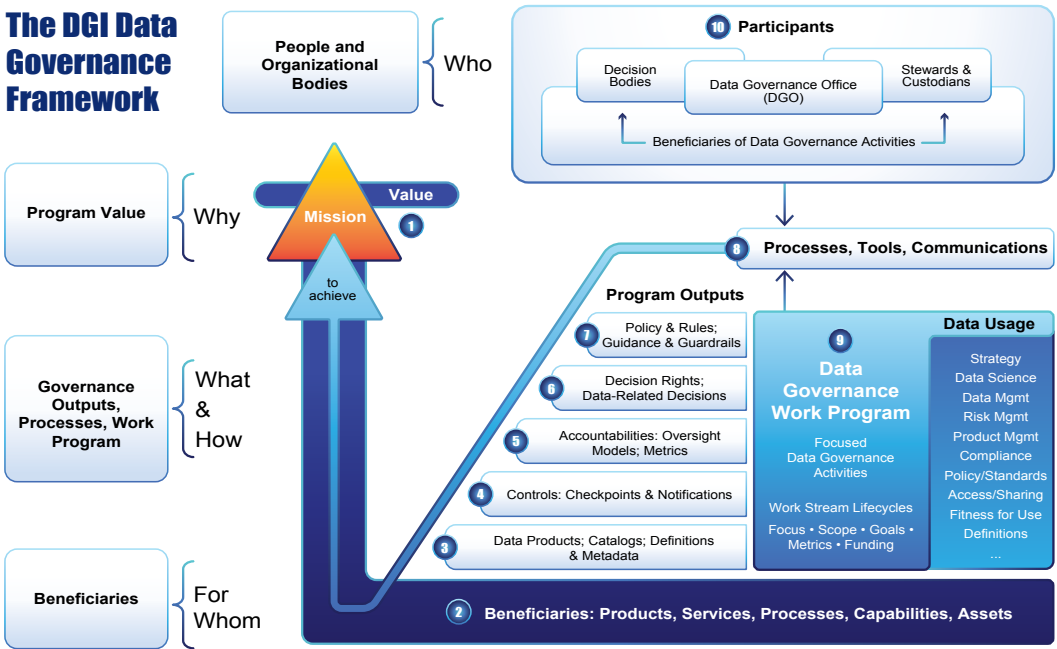
An agentic workforce operates in partnership with intelligent systems, requiring new skills that extend beyond traditional IT and software literacy.

First, an understanding of agentic AI and orchestration is critical. Employees must understand how agents are structured, how orchestration layers coordinate tasks, and how objectives translate into multi-step workflows. Workers are no longer simply users of systems; their role must evolve to that of designer and supervisor of intelligent processes. Just as operating enterprise software was once a baseline skill, knowing how to define goals, set boundaries, and interpret agent outputs is now an essential literacy. Indeed, this type of knowledge is termed 'AI literacy':

"To keep a balanced human–AI symbiosis, human decision makers must continuously update their AI literacy (e.g., how to invoke and put into practice the most recent AI developments) as well as their own competitive edge in this partnership (e.g., intuition, holistic vision, and emotional intelligence). Even though intuitive capabilities are the primary advantage of humans in decision-making, they still need to nurture analytical skills."⁶⁷

Second, data ownership and governance become a core competency. A principal thesis in our series of books about AI is that data is the most important asset. The accuracy, sovereignty, and governance of enterprise data will determine whether agents deliver quality outcomes. Employees across functions must develop skills to understand data lineage, classification, residency requirements, and privacy obligations. Sovereign cloud infrastructure becomes the foundation of trust, ensuring employees and customers know where their data resides and how it is protected. Data governance frameworks have become largely standardized.

The DGI Data Governance Framework



The DGI Data Governance Framework⁶⁸

The framework from the Data Governance Institute (above), along with the DAMA® Data Management Body of Knowledge (DAMA-DMBOK®) from DAMA International,⁶⁹ can help you define best practices, key principles, and functions of data governance for your organization.

Third, policy development becomes more important. Agents operate within specific predetermined policies, including financial controls, regulatory rules, ethical boundaries, and different risk thresholds. Employees must be able to translate legal and operational requirements into systems and policy logic. The future workforce will be able to blend traditional business domain expertise with systems thinking. The OECD AI Principles⁷⁰ and UNESCO Recommendation on the Ethics of AI⁷¹ both provide ethical frameworks that are trusted globally.

Fourth, as we covered in Chapters 3 and 13, oversight and risk management become embedded skills in the agentic enterprise. Business performance and human wellbeing are inseparable metrics. Employees must monitor agent performance, identify anomalies, evaluate potential bias, and understand escalation protocols. Performance measurement must evolve alongside these skills.

At the same time, the psychological dimension cannot be ignored. Inherent fear of AI replacing jobs must be addressed transparently. Leaders must communicate that agentic AI augments, supports, and benefits employees, rather than replacing them. Cultural adaptation will determine adoption speed. Without trust and effective change management, even the most advanced agentic system will not be successful.

Leadership for the Agentic Enterprise

Leadership in the agentic enterprise requires managing hybrid teams composed of humans and digital agents. In this context, leadership involves clarifying which decisions are autonomous, which require human review, and where shared accountability lies. Autonomy tiers must be defined, determining when agents can act independently and when escalation is mandatory.

Balancing innovation and governance has become a central theme for leaders in the agentic age. The enterprise's ability to move fast requires it to govern faster. Leaders must ensure autonomy remains bounded. Clear escalation triggers, reversible workflows, and human override mechanisms are not operational details—they are leadership decisions that must be embedded into the architecture.

Sovereignty also becomes a strategic leadership mandate. Data must remain within legal and jurisdictional boundaries. AI decisions must not expose organizations to regulatory risk. Sovereign cloud infrastructure underpins this commitment, serving as the infrastructure of trust in a hyperscale, globally connected economy.

Critically, leaders must also be in a position to determine where not to automate. Not all decisions should be delegated to agentic AI. Areas that require empathy, negotiation, and moral decision-making—e.g., complex customer complaint management or important strategic decisions—require human engagement. Leaders should shift the question from *can we automate?* to *should we?* And leaders in the agentic enterprise must draw autonomy boundaries for their teams to drive responsible AI.

Ethical Considerations for Agentic Deployments

As agentic systems expand across the enterprise, ethics must be operationalized:

- **Bias monitoring must be ongoing.** Agentic systems that act across domains like HR, finance, or customer service can heighten inequities if left unmanaged. Objective monitoring frameworks, fairness metrics, and independent audits should be included in the operating model.
- **Transparency around escalations and accountability for decision-making are essential.** When an agent makes a decision, there must be a clear record of why it acted, what data informed it, and who (humans-in-command) is ultimately responsible. Auditability and explainability are non-negotiables.
- **Sovereign compliance must be built into the deployment.** Data residency laws, regulations, privacy mandates, and other constraints must be enforced. Human trust depends on digital sovereignty. Employees and customers must know their data is governed within clear legal frameworks, without being obligated to operate under the policies of foreign jurisdictions or companies.
- **Ethical deployment requires measuring human outcomes alongside business outcomes.** Productivity uplift without burnout, confidence in AI decisions, and workforce engagement are indicators of a healthy agentic AI implementation. If automation increases output but decreases morale or trust, the model is not sustainable.

As we explored in earlier chapters, the rise of the agentic enterprise is the transition to a full agentic operating model where trusted data, orchestrated agents, and human oversight form a cohesive solution. In such a model, workflows are dynamic, cross-system, and continuous. Agents monitor, adapt, and execute 24/7. Decision cycles compress by 30–50% at minimum. Humans lead by defining objectives, managing trade-offs, governing exceptions, and providing strategic direction. The career ladder does not disappear; it evolves. Humans and agents climb it together, each in roles suited to their strengths.

The rise of the agentic enterprise, then, is not about replacing people or simply deploying smarter software. It is about rearchitecting data, workflows, governance, and culture into a new operating system for the enterprise. It is about moving from experiments to enterprise infrastructure. It is about designing organizations that learn faster, respond faster, and scale.

Success in the future enterprise will be measured by how seamlessly agentic AI is woven into its core business processes. Those who embrace this shift by balancing autonomy with oversight, speed with sovereignty, and automation with humanity will define the next generation of competitive advantage.

Executive Implications

CAIO

As autonomy increases, clear boundaries between human authority and agent execution become essential. The CAIO defines these boundaries and ensures governance evolves so speed and control remain balanced as the agentic enterprise scales.

CIO

The future agentic enterprise architecture will be hybrid by design—humans set intent, agents execute at scale. Technology leadership must define autonomy boundaries and embed governance as a native property of execution.

CFO

Competitive advantage will be measured by intelligent scale—organizations that combine workforce augmentation with disciplined governance will expand output without proportional cost growth.

CHRO

The workforce of the future is blended. Leaders must cultivate AI literacy, redefine career pathways, and ensure that automation elevates human contribution rather than eroding engagement.

CDO

Data becomes the coordination fabric between humans and agents. Sovereignty, transparency, and explainability will determine whether trust compounds—or fractures—at scale.

COO

Operational excellence evolves into orchestrated intelligence. The mandate is no longer efficiency alone, but resilient, policy-aligned execution where human oversight and machine autonomy reinforce one another.

The Agentic Enterprise: A New Operating System for Work

The rise of the agentic enterprise represents more than the adoption of a new generation of artificial intelligence. It signals the emergence of a new operating system for how organizations sense, decide, and act.

Throughout this book we have argued that agentic AI is not simply a collection of advanced tools. It is an architectural shift. Enterprise intelligence becomes distributed across agents, coordinated through orchestration layers, grounded in trusted enterprise information, and governed by humans-in-command. When these elements are designed together, organizations move beyond isolated automation toward a cohesive execution fabric capable of operating continuously, adapting in real time, and scaling intelligently.

The defining characteristic of the agentic enterprise is structured autonomy—systems that operate independently within clearly defined policies, accountability frameworks, and human oversight. In this model, governance does not slow execution; it enables it. Trust becomes the foundation that allows intelligent systems to scale safely across the enterprise.

For employees and leaders alike, this transformation elevates the human role rather than diminishing it. As agents assume responsibility for repetitive execution, humans move upward in the value chain—designing systems, setting objectives, governing outcomes, and navigating complex ethical and strategic decisions. The enterprise becomes faster not because humans disappear from the loop, but because they are positioned where their judgment matters most.

Organizations that succeed in this transition will treat agentic AI not as a series of experiments, but as infrastructure. They will build institutional memory through governed enterprise information, coordinate digital actors through orchestration layers, and embed accountability into the architecture itself. In doing so, they will create enterprises capable of learning, adapting, and executing at a scale that was previously impossible.

Endnotes

- ¹ Alexander Sukharevsky, Alexis Krivkovich, Arne Gast et al, "The agentic organization: Contours of the next paradigm for the AI era," *McKinsey & Company*, September 26, 2025, <https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/the-agentic-organization-contours-of-the-next-paradigm-for-the-ai-era>.
- ² Ibid.
- ³ Neveen Awad, Mahmood Serry, and Joe Vasquez, "How Agentic AI Is Transforming Enterprise Platforms," *Boston Consulting Group*, October 13, 2025, <https://www.bcg.com/publications/2025/how-agentic-ai-is-transforming-enterprise-platforms>.
- ⁴ Aditya Challapally, Chris Pease, Ramesh Raskar, and Pradyumna Chari, "The GenAI Divide: The State of AI in Business 2025," Project NANDA, MIT, July 2025, https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf.
- ⁵ Ibid.
- ⁶ Shaona Ghosh, Barnaby Simkin, Kyriacos Shiarlis et al. "A Safety and Security Framework for Real-World Agentic Systems," *NVIDIA*, November 27, 2025, <https://arxiv.org/pdf/2511.21990>.
- ⁷ Ted Schadler, Benjamin Nagle with Mark Moccia, Callie Smith, Ian McPherson, "Data Overview: The CIO's Role In AI Success," *Forrester*, February 4, 2026.
- ⁸ "HR Leaders to Redeploy a Quarter of Their Workforce as Agentic AI Adoption Expected to Grow 327% by 2027," *Salesforce*, May 5, 2025, <https://www.salesforce.com/news/stories/agentic-ai-impact-on-workforce-research/>.
- ⁹ "How To Lead in the Age of AI Agents," *Salesforce*, July 17, 2025, <https://www.salesforce.com/ap/blog/ai-workforce-management/>.
- ¹⁰ Ted Schadler, Benjamin Nagle with Mark Moccia, Callie Smith, Ian McPherson, "Data Overview: The CIO's Role In AI Success," *Forrester*, February 4, 2026.
- ¹¹ "EU Artificial Intelligence Act," *European Commission, the European Parliament, and the Council of the European Union*, August 1, 2024, <https://artificialintelligenceact.eu/>.
- ¹² "Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027," *Gartner*, June 25, 2025, www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027. GARTNER is a trademark of Gartner, Inc. and/or its affiliates.
- ¹³ Mark Brohan, "Gartner: AI agents will command \$15 trillion in B2B purchases by 2028," *Digital Commerce 360*, November 28, 2025, <https://www.digitalcommerce360.com/2025/11/28/gartner-ai-agents-15-trillion-in-b2b-purchases-by-2028/>. GARTNER is a trademark of Gartner, Inc. and/or its affiliates.
- ¹⁴ Rowan Curran, Leslie Joseph, Brian Hopkins, and Craig Le Clair, "Agentic AI Is the Next Competitive Frontier," *Forrester blog*, March 11, 2025, www.forrester.com/blogs/agentic-ai-is-the-next-competitive-frontier/.
- ¹⁵ Aditya Challapally, Chris Pease, Ramesh Raskar, and Pradyumna Chari, "The GenAI Divide: The State of AI in Business 2025," Project NANDA, MIT, July 2025, https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf.
- ¹⁶ "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026, Up from Less Than 5% in 2025," *Gartner*, August 26, 2025, <https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025>. GARTNER is a trademark of Gartner, Inc. and/or its affiliates.

- ¹⁷ Alexander Sukharevsky, Alexis Krivkovich, Arne Gast et al, "The agentic organization: Contours of the next paradigm for the AI era," *McKinsey & Company*, September 26, 2025, <https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/the-agentic-organization-contours-of-the-next-paradigm-for-the-ai-era>.
- ¹⁸ Harish Dunakhe, Melih Murat, and Eric Samuel, "The Rise of Agentic AI in the GCC, Riding the Next Wave of AI Revolution," *IDC*, October 2025, <https://www.etisalat.ae/content/dam/eand/assets/img/general/e-aws-whitepaper-final.pdf>.
- ¹⁹ "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026, Up from Less Than 5% in 2025," *Gartner*, August 26, 2025, <https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025>. GARTNER is a trademark of Gartner, Inc. and/or its affiliates.
- ²⁰ "Forrester Unveils Top 10 Emerging Technologies For 2025; AI Innovation Shifts From Experimentation To Business Imperative," *Forrester press release*, May 6, 2025, <https://investor.forrester.com/news-releases/news-release-details/forrester-unveils-top-10-emerging-technologies-2025-ai/>.
- ²¹ Aditya Challapally, Chris Pease, Ramesh Raskar, and Pradyumna Chari, "The GenAI Divide: The State of AI in Business 2025," Project NANDA, MIT, July 2025, https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf.
- ²² "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026, Up from Less Than 5% in 2025," *Gartner*, August 26, 2025, <https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025>. GARTNER is a trademark of Gartner, Inc. and/or its affiliates.
- ²³ Aditya Challapally, Chris Pease, Ramesh Raskar, and Pradyumna Chari, "The GenAI Divide: The State of AI in Business 2025," Project NANDA, MIT, July 2025, https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf.
- ²⁴ *Ibid.*
- ²⁵ Alexander Sukharevsky, Alexis Krivkovich, Arne Gast et al. "The agentic organization: Contours of the next paradigm for the AI era." *McKinsey & Company*, September 26, 2025. <https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/the-agentic-organization-contours-of-the-next-paradigm-for-the-ai-era>.
- ²⁶ Michele Goetz and Michael O'Grady, "AI Governance Software Spend Will See 30% CAGR from 2024 to 2030," *Forrester blog*, November 13, 2024, www.forrester.com/blogs/ai-governance-software-spend-will-see-30-cagr-from-2024-to-2030/.
- ²⁷ "Gartner Unveils Top Predictions for IT Organizations and Users in 2026 and Beyond," *Gartner*, October 21, 2025. www.gartner.com/en/newsroom/press-releases/2025-10-21-gartner-unveils-top-predictions-for-it-organizations-and-users-in-2026-and-beyond. GARTNER is a trademark of Gartner, Inc. and/or its affiliates.
- ²⁸ Carlos E. Perez, "Why Agentic AI Requires a Radically New Lifecycle," *LinkedIn*, November 7, 2025, <https://www.linkedin.com/pulse/why-agentic-ai-requires-radically-new-lifecycle-carlos-e-perez-qx7xe/>.
- ²⁹ Alexander Sukharevsky, Alexis Krivkovich, Arne Gast et al, "The agentic organization: Contours of the next paradigm for the AI era," *McKinsey & Company*, September 26, 2025. <https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/the-agentic-organization-contours-of-the-next-paradigm-for-the-ai-era>.
- ³⁰ Dan Fernandez, "The Developer's Guide to Agentic AI: The Five Stages of Agent Lifecycle Management," *DevOps.com*, October 22, 2025, <https://devops.com/the-developers-guide-to-agentic-ai-the-five-stages-of-agent-lifecycle-management/>.
- ³¹ *Ibid.*

- ³² Nancy Gohring, "From risk to reward: The dual reality of agentic AI in the enterprise," *IDC*, December 3, 2025, <https://www.idc.com/resource-center/blog/from-risk-to-reward-the-dual-reality-of-agentic-ai-in-the-enterprise/>.
- ³³ Leslie Joseph, "Announcing Our Evaluation Of The Agent Control Plane Market," *Forrester blog*, December 4, 2025, <https://www.forrester.com/blogs/announcing-our-evaluation-of-the-agent-control-plane-market/>.
- ³⁴ Dan Fernandez, "The Developer's Guide to Agentic AI: The Five Stages of Agent Lifecycle Management," *DevOps.com*, October 22, 2025, <https://devops.com/the-developers-guide-to-agentic-ai-the-five-stages-of-agent-lifecycle-management/>.
- ³⁵ Alex Singla, Alexander Sukharevsky, Bryce Hall, et al, "The state of AI in 2025: Agents, innovation, and transformation," *McKinsey & Company*, November 5, 2025, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>.
- ³⁶ Aditya Challapally, Chris Pease, Ramesh Raskar, and Pradyumna Chari, "The GenAI Divide: The State of AI in Business 2025," Project NANDA, MIT, July 2025, https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf.
- ³⁷ "The AI Divide: Why Most AI Initiatives Fail and How the Top 5% Succeed," *Marlabs*, February 17, 2026.
- ³⁸ "5 mindshifts to supercharge business growth," Global C-suite Series, *IBM*, 2025, <https://www.ibm.com/thought-leadership/institute-business-value/en-us/c-suite-study/ceo>.
- ³⁹ Alexander Sukharevsky, Dave Kerr, Klemens Hjartar et al, "Seizing the agentic AI advantage," *McKinsey & Company*, June 2025.
- ⁴⁰ Jameel Francis, "Why 85% Of Your AI Models May Fail," *Forbes*, November 15, 2024, <https://www.forbes.com/councils/forbestechcouncil/2024/11/15/why-85-of-your-ai-models-may-fail/>.
- ⁴¹ Curt Mueller, Darryl Piasecki, Marie El Hoyek et al, "How COOs maximize operational impact from gen AI and agentic AI," *McKinsey & Company*, March 20, 2025, <https://www.mckinsey.com/capabilities/operations/our-insights/how-coos-maximize-operational-impact-from-gen-ai-and-agentic-ai>.
- ⁴² Nicole Bennett, "100+ AI Statistics Shaping Business in 2026," *Vena*, January 15, 2026, <https://www.venasolutions.com/blog/ai-statistics>.
- ⁴³ Alex Singla, Alexander Sukharevsky, Lareina Yee et al, "The state of AI in early 2024: Gen AI adoption spikes and starts to generate value," *McKinsey & Company Survey*, May 30, 2024, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-2024>.
- ⁴⁴ Ted Schadler, Benjamin Nagle with Mark Moccia, Callie Smith, Ian McPherson, "Data Overview: The CIO's Role In AI Success," *Forrester*, February 4, 2026.
- ⁴⁵ Alex Singla, Alexander Sukharevsky, Lareina Yee et al. "The state of AI in early 2024: Gen AI adoption spikes and starts to generate value," *McKinsey & Company Survey*, May 30, 2024, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-2024>.
- ⁴⁶ Alex Singla, Alexander Sukharevsky, Bryce Hall et al, "The state of AI in 2025: Agents, innovation, and transformation," *McKinsey & Company*, November 5, 2025, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>.
- ⁴⁷ Tom Coshov, Arnold Gao, Lawrence Pingree et al, "Top Strategic Technology Trends for 2025: Agentic AI," *Gartner*, October 21, 2024, <https://www.gartner.com/en/documents/5850847>. GARTNER is a trademark of Gartner, Inc. and/or its affiliates.
- ⁴⁸ Emily Zhang, "Why 95% of GenAI projects fail—and why the 5% that survive matter," *Trullion Blog*, September 8, 2025, <https://trullion.com/blog/why-95-of-ai-projects-fail-and-why-the-5-that-survive-matter/>.

- ⁴⁹ Jessica Apotheker, Vinciane Beauchene, Nicolas de Bellefonds et al, "The Widening AI Value Gap," *Boston Consulting Group*, September 2025, <https://media-publications.bcg.com/The-Widening-AI-Value-Gap-October-2025.pdf>.
- ⁵⁰ Ibid.
- ⁵¹ Ibid.
- ⁵² Ibid.
- ⁵³ Ibid.
- ⁵⁴ Mark Brohan, "Gartner: AI agents will command \$15 trillion in B2B purchases by 2028," *Digital Commerce 360*, November 28, 2025, <https://www.digitalcommerce360.com/2025/11/28/gartner-ai-agents-15-trillion-in-b2b-purchases-by-2028/>. GARTNER is a trademark of Gartner, Inc. and/or its affiliates.
- ⁵⁵ "Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027." *Gartner*, June 25, 2025. www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027. GARTNER is a trademark of Gartner, Inc. and/or its affiliates.
- ⁵⁶ Ibid.
- ⁵⁷ S. Ransbotham, D. Kiron, S. Khodabandeh, S. Iyer, and A. Das, "The Emerging Agentic Enterprise: How Leaders Must Navigate a New Age of AI," *MIT Sloan Management Review and Boston Consulting Group*, November 2025.
- ⁵⁸ "Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions," *OECD*, September 18, 2025, https://www.oecd.org/en/publications/governing-with-artificial-intelligence_795de142-en/full-report/ai-in-public-service-design-and-delivery_09704c1a.html.
- ⁵⁹ Ted Schadler, Benjamin Nagle with Mark Moccia, Callie Smith, Ian McPherson, "Data Overview: The CIO's Role In AI Success," *Forrester*, February 4, 2026.
- ⁶⁰ Josh Taylor, "Facebook outage: what went wrong and why did it take so long to fix after social platform went down?" *The Guardian*, October 5, 2021, <https://www.theguardian.com/technology/2021/oct/05/facebook-outage-what-went-wrong-and-why-did-it-take-so-long-to-fix>.
- ⁶¹ "Artificial Intelligence Risk Management Framework (AI RMF 1.0)," *National Institute of Standards and Technology – U.S. Department of Commerce*, January 2023, <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>.
- ⁶² Ibid.
- ⁶³ "EU Artificial Intelligence Act," *European Commission, the European Parliament, and the Council of the European Union*, August 1, 2024, <https://artificialintelligenceact.eu/>.
- ⁶⁴ Inioluwa Deborah Raji, Andrew Smart, Rebecca N. White et al, "Closing the AI accountability gap: defining an end-to-end framework for internal algorithmic auditing," *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, January 27, 2020, <https://dl.acm.org/doi/pdf/10.1145/3351095.3372873>.
- ⁶⁵ Sebastian Raisch and Sebastian Krakowski, "Artificial Intelligence and Management: The Automation–Augmentation Paradox." *Academy of Management Review*, 2021, 46(1), 192–210, https://www.researchgate.net/profile/Sebastian-Raisch/publication/348483816_Artificial_Intelligence_and_Management_The_Automation-Augmentation_Paradox/links/6755c5e1ad10b614ef38f8bd/Artificial-Intelligence-and-Management-The-Automation-Augmentation-Paradox.pdf.
- ⁶⁶ Ibid.
- ⁶⁷ M. H. Jarrahi, "Artificial Intelligence and the Future of Work: Human–AI Symbiosis in Organizational Decision Making," *Business Horizons*, 2018, 61(4), 577–586.

⁶⁸ DGI Data Governance Framework, *The Data Governance Institute*, <https://datagovernance.com/the-dgi-data-governance-framework/>.

⁶⁹ The DAMA® Data Management Body of Knowledge (DAMA-DMBOK®), *DAMA International*, <https://dama.org/learning-resources/dama-data-management-body-of-knowledge-dmbok/>.

⁷⁰ "AI principles," *OECD*, 2024, <https://www.oecd.org/en/topics/sub-issues/ai-principles.html>.

⁷¹ "Ethics of Artificial Intelligence," *UNESCO*, November 2021, <https://www.unesco.org/en/artificial-intelligence/recommendation-ethics>.

Glossary of Terms

Agent Architects: People who design agentic systems within the enterprise. They define workflow logic, autonomy tiers, escalation triggers, and integration points to ensure agents operate safely and coherently across systems of record. Working within established governance and security guardrails, they translate business objectives and policy requirements into scalable, policy-aligned orchestration designs.

Agentic Artificial Intelligence (AI): The framework of artificial intelligence systems designed to function as autonomous agents. Unlike models that simply respond to a prompt, an agent can perceive its environment, create a multi-step plan, make independent decisions, and use tools to actively work toward a specific goal. Agentic AI describes the overarching paradigm within which individual “AI agents” operate.

Agentic AI Genome: The new enterprise operating model for agentic AI, defining how governed enterprise information, specialized agents, orchestration layers, and human-in-command oversight are structured into a scalable, policy-aligned execution system. It provides a repeatable architectural pattern that enables autonomous workflows to operate across domains while preserving accountability, sovereignty, and lifecycle control.

Agentic AI Genome Map: The visual architectural model that illustrates how the agentic enterprise is structured across layered components—enterprise content and systems (the substrate), domain agents, orchestration and control layers, and human-in-command oversight. It serves as a blueprint for scaling autonomous workflows consistently across business domains while embedding governance, policy enforcement, auditability, and sovereignty by design.

Agent Control Plane: The governance and coordination layer of the agentic enterprise that defines, monitors, constrains, and audits how agents operate across workflows. It sits above individual agents and below human strategic intent, ensuring that autonomous execution remains aligned with enterprise policy, risk thresholds, data sovereignty requirements, and human-in-command oversight. Where the agent layer performs tasks and executes decisions, the control plane defines the rules of engagement.

Agentic Enterprise: An organization that embeds autonomous, policy-governed agents into its core operating model, enabling workflows to sense, decide, and act across systems with human-in-command oversight. Grounded in trusted enterprise information and coordinated through orchestration and control layers, the agentic enterprise moves beyond advisory AI to outcome-driven execution—scaling velocity, resilience, and accountability together through governed, sovereign architecture.

AI Agent: A specific software system designed for a defined task that can autonomously perceive its environment, plan and reason about tasks, and take actions to achieve specific goals. It operates with a degree of independence (but usually within human-set constraints), learns or adapts over time, uses tools or external data sources when needed, and supports decision-making with minimal direct supervision. AI agents are the building blocks of the framework of agentic AI.

AI-Driven Analytics: The application of artificial intelligence and machine learning to automate data collection, preparation, and analysis. By uncovering patterns and generating insights in real or near real time, it helps organizations predict outcomes, identify trends, and make faster, more informed decisions.

AI Governance Manager or Supervisor: A designated leadership role responsible for defining, monitoring, and enforcing the policies, controls, and compliance standards that govern AI and agentic systems across the enterprise. This role oversees the development of autonomy tiers, escalation protocols, data access controls, model validation practices, and audit mechanisms to ensure that intelligent systems operate within defined legal, regulatory, ethical, and risk boundaries.

AI Ops Engineer: Responsible for the real-time monitoring, reliability, and operational performance of AI and agentic systems in production. This role oversees telemetry, logging, drift detection, anomaly identification, and incident response to ensure agents operate within defined performance and policy thresholds. Working closely with governance and architecture teams, the AI Ops Engineer investigates irregular behaviors, triggers escalation protocols, and supports remediation or rollback when necessary.

AI-Powered Assistant (also called copilots): Systems designed to augment human work by providing recommendations, generating content, summarizing information, or supporting task completion within existing workflows. Unlike agentic systems, AI-powered assistants do not independently execute end-to-end processes; they operate under direct human initiation and oversight, with the human retaining final decision authority.

AI-Ready Platform/Stack: A technology environment built to support the development, deployment, and scaling of artificial intelligence applications. It combines scalable infrastructure, strong data governance, and modular tools—such as APIs, model management, and integration pipelines—to make AI projects faster, more reliable, and easier to move from experimentation to production.

Analytics: The systematic process of collecting, processing, and interpreting data in order to identify patterns, trends, and insights. In computing, analytics is used to support decision-making, optimize performance, and predict future outcomes through techniques such as statistical analysis, data visualization, and machine learning.

API (Application Programming Interface): A set of software rules and protocols that enable two applications to communicate with each other. APIs provide a structured way for AI systems to connect programmatically with external models, datasets, or other software components.

Application Security as a Service (ASaaS): A cloud-based model for continuously assessing and managing application security across its lifecycle. It provides centralized scanning, policy enforcement, and actionable insights to identify and remediate vulnerabilities before and after deployment. In an agentic AI context, ASaaS acts as a control layer that enables continuous monitoring and supports increasingly automated, policy-driven security workflows.

Artificial Intelligence (AI): Technology that empowers machines or software to perform tasks normally requiring human intelligence—things like learning, reasoning, perception, problem-solving, decision-making, natural language understanding, and recognizing patterns. AI systems mimic or simulate cognitive functions of the human mind by using algorithms, large datasets, and computational power.

Audit (in computing): A systematic examination and evaluation of systems, processes, or data in computing to ensure accuracy, security, compliance, and proper operation. Audits may review logs, access controls, configurations, and policies to detect irregularities, confirm adherence to standards, and identify areas for improvement.

Audit Trails: Chronological records that track system activities, including user actions, data changes, and access events. They provide transparency, support regulatory compliance, and help with security investigations by showing who did what, when, and how within a system.

Auditable Access: The capability of a system to log and review who accessed resources, when, and how. It provides accountability and compliance by ensuring access activity can be verified and traced if needed.

Automation: The use of technology to perform tasks with minimal or no human intervention. In computing, automation streamlines repetitive or rule-based processes—such as software deployment, system monitoring, or data processing—to improve efficiency, consistency, and reliability.

Bots: Software programs designed to automate tasks, often by simulating human activity. In computing, bots can perform a wide range of functions at high speed and scale, from helpful activities such as customer support chatbots, search engine indexing, and workflow automation, to malicious uses like spamming, credential stuffing, or spreading malware.

Business Intelligence (BI): The technologies, tools, and practices that collect, integrate, and analyze business data to support decision-making. BI platforms transform raw data into dashboards, reports, and visualizations, helping organizations identify trends, measure performance, and make more informed strategic choices.

Business Process Management (BPM): Refers to aligning processes with an organization's strategic objectives, designing and implementing process-centric tools or architectures, and determining measurement systems for effective process management.

CAC (Customer Acquisition Cost): A metric that measures the total cost required to acquire a new customer. It includes expenses related to marketing, sales, advertising, and onboarding activities. CAC helps organizations evaluate the efficiency of their growth strategies and the sustainability of customer acquisition efforts.

Change Management: A discipline that focuses on the human and organizational side of change, ensuring that new systems, processes, or technologies are adopted successfully. While configuration management deals with technical consistency, change management addresses communication, training, and stakeholder alignment to minimize resistance and maximize value from change initiatives.

Chief AI Officer (CAIO): The executive leader responsible for enterprise-wide AI strategy, governance, and value realization. As the primary architect of AI governance, the CAIO works in concert with the CIO, CHRO, COO, and CISO to define autonomy boundaries, policy frameworks, security controls, and risk guardrails that business and domain owners must enforce. The role ensures that AI and agentic systems align with business objectives while meeting regulatory, ethical, and sovereignty requirements.

Cloud/The Cloud: A model of computing that delivers on-demand access to shared resources—such as servers, storage, databases, networking, software, and analytics—over the internet. The cloud allows users and organizations to scale resources quickly, pay only for what they use, and access services without maintaining physical hardware or infrastructure onsite.

Compliance (in computing): The practice of ensuring that systems, processes, and data management meet established laws, regulations, standards, and internal policies. In tech, compliance often covers areas like data privacy (e.g., GDPR, CCPA), cybersecurity, accessibility, and industry-specific rules, helping organizations reduce risk, maintain trust, and avoid legal or financial penalties. (See also: Data Security; Cybersecurity; Data Privacy).

Computer Vision: A field of AI that enables computers and systems to derive meaningful information from digital images, videos, and other visual inputs, allowing them to take action or make recommendations. It interprets visual data by identifying and analyzing objects, scenes, and actions.

Configuration Management: The process of systematically handling changes to a system to maintain its integrity, consistency, and traceability throughout its lifecycle.

Content Lifecycle Management (CLM): The combination of document management, records management, workflow, archiving, and imaging into a fully integrated solution to effectively manage the lifecycle of content, from creation through to archiving and eventual deletion.

Contextual Intelligence: The ability of systems, organizations, or individuals to interpret data, events, or behaviors within their surrounding context and act appropriately. In technology, it refers to AI's use of real-time signals—such as location, behavior, preferences, or timing—to deliver more relevant insights, recommendations, and actions.

Control Plane: See Agent Control Plane.

CPQ (Configure, Price, Quote): Software systems and processes that help organizations configure products or services, calculate accurate pricing, and generate sales quotes. CPQ streamlines complex sales workflows by applying pricing rules, product options, and discount policies automatically. It improves sales efficiency, pricing accuracy, and customer responsiveness.

Customer Relationship Management (CRM): A strategy and set of software tools that help businesses manage interactions with current and potential customers. CRM systems centralize customer data, track sales and communications, and support marketing, service, and relationship-building to improve retention and drive growth.

Cybersecurity: The practice of protecting systems, networks, software, and data from digital attacks, unauthorized access, or damage. It encompasses technologies, processes, and policies that safeguard confidentiality, integrity, and availability, helping organizations defend against threats like malware, phishing, ransomware, and insider risks.

Data Privacy: The discipline of managing and protecting personal information to ensure it is collected, stored, and used in ways that respect individuals' rights and expectations. Data privacy focuses on transparency, consent, and legal compliance, ensuring that organizations handle sensitive data responsibly while maintaining user trust. (See also: Data Security; Cybersecurity; Compliance).

Data Residency: The physical or geographic location where an organization's data is stored and processed. Often driven by legal, regulatory, or business requirements, data residency ensures that information remains within specific jurisdictions, which can impact compliance, security, and performance.

Data Security: The set of practices, technologies, and policies used to protect digital information from unauthorized access, corruption, or loss. It encompasses measures such as encryption, access controls, backups, and monitoring to safeguard the confidentiality, integrity, and availability of data across its lifecycle. (See also: Cybersecurity; Compliance; Data Privacy).

Data Silos: Isolated collections of data that are controlled by one department, system, or platform and are not easily accessible to others within an organization. Data silos limit collaboration, reduce visibility, and can create inefficiencies or inconsistencies, making it harder to gain a unified view of information across the business.

Data Sovereignty: Data sovereignty is the principle that digital data is subject to the laws and governance of the country or jurisdiction where it is collected, stored, or processed. It requires organizations to manage data in compliance with local regulations, including privacy, security, and data residency requirements. In cloud environments, sovereignty can be challenged because third-party providers may be subject to foreign legal authorities, potentially allowing external governments to access data even when it is stored domestically. Maintaining data sovereignty therefore requires architectural controls and governance policies that ensure sensitive information remains protected under the appropriate legal jurisdiction.

Deep Learning: A specialized branch of machine learning that relies on deep neural networks—multi-layered structures of interconnected “neurons” whose weights and parameters can be adjusted through training. This approach excels at extracting patterns and insights from unstructured data such as images, text, audio, and video, making it the backbone of many modern AI applications like image recognition, language translation, and speech processing.

Digital Workforce: A collection of automated software systems, known as “digital workers” (such as AI agents, bots, and virtual assistants), that perform tasks traditionally done by humans.

Digital Twin: A virtual representation of a physical asset, system, or process that is continuously updated with real-world data to reflect its current state, behavior, and performance. In an enterprise context, digital twins are used to monitor operations, run simulations, and predict outcomes under different conditions.

Edge Computing: A distributed computing architecture in which processing power, storage, and data analysis are located nearer to where data is generated—on devices, gateways, or local “edge” servers—rather than centralized in remote cloud data centers. This approach improves response time, reduces latency and bandwidth usage, and enables real-time or near-real-time applications, particularly in IoT, autonomous systems, and environments where speed or local decision-making matters.

EDI (Electronic Data Interchange): A standardized method for exchanging business documents—such as purchase orders, invoices, and shipping notices—electronically between organizations, eliminating manual data entry and improving speed, accuracy, and consistency in transactions.

Encryption: The process of converting data into a coded format using algorithms and cryptographic keys to prevent unauthorized access. In computing, encryption ensures that only authorized parties with the correct key can decrypt and read the information, protecting sensitive data during storage or transmission.

Enterprise AI (EAI): The disciplined application of artificial intelligence within an organization to solve real business problems, improve decision-making, and automate work securely and at scale. It is not a separate category of intelligence, but the governed deployment of existing AI capabilities—including machine learning, natural language processing, computer vision, automation, and generative AI—within a structured data and compliance environment.

Enterprise Content Management (ECM): A platform that stores, manages, and delivers enterprise-level content. This includes documents, images, videos, and other forms of content that are important to an organization. An ECM platform should seamlessly integrate with crucial enterprise applications and systems (such as enterprise resource planning, customer relationship management, human capital management, and supply chain management solutions) to accelerate business processes and leverage the data they generate.

Enterprise Information Management (EIM): Solutions that manage the creation, capture, use, and eventual lifecycle of structured and unstructured information. They are designed to help organizations extract value from their information, secure that information, and meet the growing list of compliance requirements.

Enterprise Resource Planning (ERP): An integrated software system that manages core business processes—such as finance, supply chain, manufacturing, human resources, and customer relations—within a unified platform.

Enterprise Service Management (ESM): The practice of applying service management principles and platforms—traditionally used in IT—to other business functions such as HR, finance, facilities, and customer support. It standardizes how services are requested, delivered, and managed across the enterprise. ESM improves operational efficiency, transparency, and consistency in service delivery.

Fallout Rates: The percentage of automated transactions or workflow steps that fail to complete autonomously and require manual intervention, correction, or escalation. In agentic systems, fallout rates measure the reliability of end-to-end intelligent workflows and the frequency of exceptions that fall outside automated decision boundaries.

General Data Protection Regulation (GDPR): A comprehensive data privacy and protection law enacted by the European Union in 2018. GDPR governs how organizations collect, process, store, and share personal data, emphasizing transparency, user consent, and individual rights, with significant penalties for non-compliance.

Generative AI (GenAI): AI systems that generate content using machine learning models. GenAI systems such as ChatGPT, Claude, Gemini, and DeepSeek are trained on public data sources like websites, news, Reddit, and Wikipedia. While public GenAI chatbots are useful for generating general insights, they are limited to general-purpose tasks. This is because they lack access to the private, real-time, and enterprise-centric data required for specific business use cases.

Governance: The framework of policies, roles, processes, and standards that guide how an organization manages its digital assets, technologies, and data. This same structure guides how AI and agentic systems are designed, deployed, monitored, and retired. Governance ensures that intelligent systems operate within defined legal, regulatory, ethical, security, and risk boundaries while remaining aligned to enterprise objectives.

HIPAA (Health Insurance Portability and Accountability Act): A U.S. law that establishes national standards for protecting sensitive patient health information. It requires organizations handling medical data to implement safeguards for privacy, security, and controlled access. HIPAA compliance ensures that personal health information is protected from unauthorized use or disclosure.

Human in Command (HIC): A model for AI where a human defines the strategy, goals, and ethical boundaries, but the AI acts autonomously within those boundaries. The human acts as a commander setting the rules, not a worker checking every step.

Human in the Loop (HITL): A model where a human is an active part of the decision-making process. The AI handles data processing but flags high-stakes or uncertain cases for a human to review, approve, or correct before action is taken.

Human Oversight Owner: The designated accountable leader for a specific agent or agentic workflow, responsible for ensuring its decisions and actions remain aligned with business objectives, policy constraints, and ethical standards. Operating within the governance framework established by the Chief AI Officer (CAIO), this role monitors performance, reviews escalations, validates exceptions, and retains authority to intervene, override, or suspend autonomous execution when required.

HSE (Health, Safety, and Environment): The policies, practices, and regulatory frameworks that ensure workplace health, employee safety, and environmental protection. Organizations implement HSE programs to manage operational risks, prevent accidents, and comply with safety and environmental regulations.

Hybrid Cloud: A cloud architecture that integrates on-premises infrastructure with public and/or private cloud environments, allowing workloads and data to operate across multiple environments as a unified system. In an agentic enterprise, hybrid deployment enables organizations to balance scalability and innovation with regulatory, sovereignty, and security requirements by keeping sensitive data or high-risk workflows within controlled environments while leveraging cloud elasticity for broader processing needs.

Hyperscaler: A large cloud service provider—such as Amazon Web Services (AWS), Microsoft Azure, or Google Cloud—that delivers massive-scale computing, storage, and networking services across global data centers. Hyperscalers are known for their ability to scale resources up or down instantly, support multi-tenant workloads, and provide the backbone for cloud-native, AI, and data-intensive applications.

I2P (Invisible Internet Project): A decentralized network designed to enable secure, anonymous communication over the internet. It routes traffic through multiple encrypted layers across volunteer-operated nodes to conceal the origin and destination of data. I2P is often used to protect user privacy and resist network surveillance.

Identity and Access Management (IAM): A security framework that ensures the right individuals have the appropriate access to the right resources at the right time. It manages digital identities, authentication, and permissions across systems and applications to protect sensitive data and maintain compliance.

In Silico: Experiments, simulations, or analyses performed using computer models rather than physical or real-world testing. The term is commonly used in scientific research and AI development to evaluate hypotheses, train models, or test scenarios within a digital environment. In enterprise AI contexts, in silico methods allow organizations to simulate workflows, stress-test systems, and evaluate outcomes before deployment in production environments. In silico methods often leverage digital twins—virtual representations of physical assets, processes, or entire organizational workflows—especially in medicine and engineering.

Intelligence Layer: The analytical and decision-making tier within a technology stack that transforms raw data into actionable insights. Often powered by AI, machine learning, or advanced analytics, the intelligence layer sits above data storage and processing systems, enabling personalization, predictions, and automations that drive smarter business outcomes.

Internet of Things (IoT): A network of physical devices—such as sensors, appliances, vehicles, and machinery—that are connected to the internet and can collect and share telemetry data, enabling automation, monitoring, and smarter decision-making in real time.

ITAR (International Traffic in Arms Regulations): A U.S. regulatory framework that controls the export, access, and transfer of defence-related technologies, data, and services. It requires organizations handling controlled defense information to implement strict security and access controls. ITAR compliance ensures that sensitive military technologies are protected from unauthorized disclosure or foreign access.

JSON (JavaScript Object Notation): A lightweight data format used to structure and exchange information between systems. It organizes data in simple key-value pairs and nested objects that are easy for both humans and machines to read and process. JSON is widely used in APIs, web services, and AI systems to transmit structured data.

KYC (Know Your Customer/Client): A mandatory, regulated process used by financial institutions to verify a customer's identity, assess risk, and prevent illegal activities like money laundering, fraud, and terrorism financing.

Key Performance Indicators (KPIs): Quantifiable metrics used to evaluate how effectively an organization, team, or process is achieving its objectives. KPIs track progress toward strategic goals, guide decision-making, and can range from financial measures like revenue growth to operational ones like customer retention or system uptime.

Kubernetes: An open-source platform for automating the deployment, scaling, and management of containerized applications. Originally developed by Google, Kubernetes orchestrates clusters of containers—handling scheduling, load balancing, and failover—so applications run reliably and efficiently across cloud, on-premises, or hybrid environments.

Laboratory Information Management System (LIMS): Software used to manage laboratory workflows, samples, test results, and associated data. It helps laboratories track samples from collection through analysis while ensuring data integrity, traceability, and regulatory compliance.

Large Language Model (LLM): A type of foundational model built with deep learning that is pre-trained on enormous text corpora using self-supervised methods. These models process inputs in the form of “tokens” (pieces of words or characters), learn the statistical relationships among them, and use those relationships to understand, generate, summarize, or transform human-language text.

Least-Privilege Access: A security principle that limits users, systems, or agents to only the permissions necessary to perform their specific tasks. By restricting access to the minimum required data and functions, organizations reduce the risk of unauthorized actions, data exposure, or misuse. In agentic systems, it ensures agents operate within clearly defined boundaries.

Legacy Systems: Outdated or older technology systems, software, or infrastructure that remain in use despite being superseded by newer alternatives. Legacy platforms may still support critical business operations but often present challenges such as limited compatibility, higher maintenance costs, security vulnerabilities, and difficulty integrating with modern solutions.

Machine Learning (ML): A field of artificial intelligence that focuses on developing algorithms and models that enable systems to learn from data and improve their performance over time without being explicitly programmed. ML is used to identify patterns, make predictions, and support decision-making in applications such as recommendation engines, fraud detection, image recognition, and natural language processing.

Manufacturing Execution System (MES): Software that manages and monitors production processes on the factory floor. It tracks materials, equipment, and work-in-progress to ensure manufacturing operations follow defined procedures and quality standards. MES provides real-time visibility into production performance and helps improve efficiency, traceability, and compliance.

Metadata: Any set of terms, words, symbols, and numbers embedded within a document to allow records management functions such as classification, search, historical tracking (date created, modified, retrieved), and user identification (authors and editors of each refinement). This structured information ensures the document's characteristics remain traceable and manageable throughout its lifecycle.

Middleware Layer: Software that acts as a bridge between operating systems, databases, and applications, enabling them to communicate and share data efficiently. The middleware layer provides common services such as messaging, authentication, API management, and transaction processing—simplifying integration and interoperability across complex systems.

Model Context Protocol (MCP): A standardized framework that allows AI models and agents to securely access external data sources, tools, and services. By providing structured connections to enterprise systems, MCP enables agents to retrieve real-time context, perform actions, and interact with business workflows while maintaining governance, permissions, and auditability.

Modernization (in software): The broader process of updating legacy systems, applications, or infrastructure to take advantage of modern technologies, architectures, and practices. While migration often focuses on relocating existing systems, modernization may involve refactoring, replatforming, or redesigning to improve scalability, agility, and long-term business value.

Monolithic Architecture: A traditional software design where all components of an application—such as the user interface, business logic, and data management—are tightly integrated and deployed as a single unit. While simpler to build initially, monolithic systems can be harder to scale, update, or adapt compared to composable alternatives.

MTTD (Mean Time to Detect): A metric that measures the average time it takes to identify a system failure, anomaly, or security incident after it occurs. It is used in IT operations, cybersecurity, and reliability engineering to evaluate how quickly issues are discovered. Lower MTTD indicates faster detection and improved operational visibility.

MTTR (Mean Time to Repair): A metric that measures the average time required to diagnose, fix, and restore a system or service after a failure. It is commonly used in IT operations and reliability engineering to evaluate incident response and recovery performance. Lower MTTR indicates faster resolution of issues and improved operational resilience.

Multi-Cloud: A cloud strategy in which an organization uses services from two or more cloud providers simultaneously. This approach helps reduce vendor lock-in, increase resilience, optimize performance, and meet regulatory or geographic requirements by distributing workloads across multiple platforms such as AWS, Microsoft Azure, and Google Cloud.

Multi-Region Model: A cloud computing architecture in which applications and data are deployed across multiple geographic regions offered by a cloud provider. This model improves availability, performance, and disaster recovery by distributing workloads closer to end users and ensuring redundancy if one region experiences outages or latency issues.

Named Entity Recognition (NER): A natural language processing technique that identifies and classifies key entities in text, such as people, organizations, locations, dates, or product names. It helps AI systems extract structured information from unstructured language. NER is commonly used in search, document analysis, and enterprise AI workflows to support information retrieval and automation.

Natural Language Processing (NLP): A field of artificial intelligence that enables computers to understand, interpret, and generate human language. NLP combines linguistics, machine learning, and computational techniques to support applications such as chatbots, translation, sentiment analysis, and voice recognition.

Natural Language Understanding (NLU): A subset of artificial intelligence and Natural Language Processing (NLP) that enables machines to comprehend the intent, context, and meaning behind human language. It allows AI models to recognize intent, context, sentiment, and relationships within spoken or written text, transforming unstructured language into structured representations that machines can act upon.

Optical Character Recognition (OCR): Technology that converts printed or handwritten text in scanned documents or images into machine-readable digital text, enabling search, editing, and automated processing.

Orchestration Layer: A management layer in computing that automates the coordination, scheduling, and execution of complex tasks across multiple systems, applications, or services. The orchestration layer ensures that components work together seamlessly by handling workflows, resource allocation, scaling, and dependencies. It is commonly used in cloud environments, containerized applications, and microservices to streamline operations and reduce manual intervention.

Orchestrator: The coordination layer within an agentic system that manages how multiple agents, services, and workflows interact to complete a task or achieve a defined objective. It directs the sequence of actions, routes information between agents and systems of record, applies policy constraints, and triggers escalation when predefined thresholds are reached. In the agentic enterprise, the orchestrator ensures that specialized agents operate as part of a coherent, governed workflow rather than as isolated components.

P2P (Peer-to-Peer): A decentralized networking model in which computers (peers) communicate and share resources directly with one another without relying on a central server. Each participant can act as both a client and a provider of data or services.

Permissions: Management of who can access a computer or network. The Access Control List (ACL) is the set of data associated with a file, directory, or other resource that defines the permissions that users, groups, processes, or devices have for accessing it.

Personal Information Protection and Electronic Documents Act (PIPEDA): Canada's federal privacy law governing how private-sector organizations collect, use, and disclose personal information in the course of commercial activities. PIPEDA grants individuals rights to access and correct their data, requires organizations to obtain meaningful consent, and mandates safeguards to protect personal information.

Personally Identifiable Information (PII): Any data that can be used to identify an individual, either on its own or when combined with other information. Examples include names, addresses, phone numbers, email addresses, Social Security or passport numbers, and financial or health records. In computing and data privacy, protecting PII is critical to complying with regulations and safeguarding individuals from identity theft or misuse.

PPM (Project Portfolio Management): The practice of selecting, prioritizing, and managing a portfolio of projects to align with an organization's strategic objectives. It provides visibility into project performance, resource allocation, and risk across the portfolio. PPM helps organizations ensure that investments in initiatives deliver maximum value and remain aligned with business priorities.

Privacy: The protection and proper handling of user data to ensure individuals maintain control over how their personal information is collected, used, shared, and stored in computing environments. Privacy safeguards prevent unauthorized access or misuse of data and are guided by legal, ethical, and regulatory standards. (See also: Data Privacy).

Privacy-by-Design: A framework that embeds privacy and data protection principles directly into the design and operation of technologies, processes, and systems. It emphasizes proactive measures—such as minimizing data use, safeguarding by default, and ensuring transparency—so privacy is not an afterthought but a core design principle.

Private Cloud: A cloud computing environment dedicated to a single organization, offering exclusive access to infrastructure, resources, and services. Private clouds can be hosted on-premises or by a third-party provider, and are designed to deliver greater control, security, and customization compared to public cloud environments. (See also: Public Cloud, Hybrid Cloud).

Process Management: The automation of business processes using a rule-based expert system that invokes the appropriate tools and supplies the necessary information, checklists, examples, and status reports to the user.

Prompt (and Interaction) Engineer: Responsible for designing how agents interpret context, process human input, and generate responses in a consistent, reliable, and policy-aligned manner. This role defines structured prompts, contextual frameworks, and interaction patterns that ensure agent behavior remains predictable across varied use cases.

Public Cloud: A cloud computing model in which infrastructure, resources, and services are owned and operated by a third-party provider and delivered over the internet. Public cloud environments are shared among multiple organizations (tenants) but keep data and workloads separated. They offer scalability, flexibility, and cost efficiency, with common examples including Amazon Web Services (AWS), Microsoft Azure, and Google Cloud.

RPA (Robotic Process Automation): A technology that uses software bots to automate repetitive, rule-based tasks typically performed by humans. These bots interact with applications and systems to perform actions such as entering data, processing transactions, and triggering workflows. RPA improves efficiency and reduces manual effort by automating routine business processes.

RTO (Recovery Time Objective): The maximum acceptable amount of time a system, application, or service can remain unavailable after a disruption before it must be restored. It is used in disaster recovery and business continuity planning to define recovery priorities and response strategies. Shorter RTO targets indicate a greater need for rapid system restoration to minimize operational impact.

Real-Time Analytics: The process of collecting, processing, and analyzing data immediately as it is generated, allowing organizations to gain insights and make decisions without delay. In computing, real-time analytics supports use cases such as fraud detection, personalized recommendations, system monitoring, and live performance tracking.

Red-Teaming (Agentic AI): The structured practice of challenge-driven/attack-simulation testing of agentic AI systems to identify security vulnerabilities, governance gaps, bias risks, and policy weaknesses before deployment at scale. In the agentic enterprise, it extends beyond model testing to include workflow integrity, autonomy boundaries, escalation controls, and data sovereignty—ensuring that autonomous execution remains aligned with human oversight and enterprise policy.

Repository: A storage location, often managed with tools like GitHub or GitLab, where a codebase and its history of changes are kept. While the codebase refers to the actual code itself, the repository also tracks revisions, branches, and contributions, enabling teams to manage and collaborate on the code effectively.

Retrieval-Augmented Generation (RAG): An AI architecture that enhances generative models by retrieving relevant information from external data sources—such as enterprise documents, databases, or knowledge repositories—before generating a response. In the agentic enterprise, RAG enables agents to access governed enterprise content and systems of record to support decisions, actions, and responses. This approach helps reduce hallucinations while aligning AI outputs with trusted organizational data.

Rights and Permissions: Identifies the circumstances under which a particular asset may be used. For instance, it indicates who legally owns the asset, in what mediums it may be used (web, print, T5), and the financial liabilities incurred to include the asset.

Rules Engine: A system that makes decisions and applies logic to workflows based on predefined business rules.

SCADA (Supervisory Control and Data Acquisition): An industrial control system used to monitor and manage equipment, infrastructure, and industrial processes in real time. It collects data from sensors and devices across facilities such as power plants, manufacturing lines, and utilities. SCADA systems enable operators to supervise operations, analyze performance, and respond to issues quickly.

SLA (Service Level Agreement): A formal, binding contract between a service provider and a customer that defines expected service levels, such as uptime, performance, and responsibilities.

Scalability: The ability of a system, application, or infrastructure to handle increasing workloads or demand by adding resources such as processing power, memory, or storage. In computing, scalability ensures consistent performance and reliability as usage grows and can apply to scaling up (vertical scaling) or scaling out (horizontal scaling).

Sentiment Analysis Tools: Software applications that use natural language processing, machine learning, or statistical methods to identify and categorize opinions or emotions expressed in text, speech, or other data. These tools help organizations determine whether sentiment is positive, negative, or neutral, and are commonly used in areas such as customer feedback, social media monitoring, and market research.

Service-Oriented Architecture (SOA): A software design approach in which applications are built as a collection of interoperable services that communicate over a network. Each service performs a specific business function and can be reused across multiple systems and workflows. SOA enables organizations to integrate disparate systems, improve flexibility, and support scalable, modular enterprise architectures.

Sovereign Zone: A controlled computing environment where data, AI models, and agentic workflows operate under specific jurisdictional, regulatory, and security requirements. It ensures that sensitive information remains within defined geographic or legal boundaries while enforcing strict governance, access controls, and compliance policies. In the agentic enterprise, Sovereign Zones allow autonomous systems to operate safely while preserving data residency, regulatory compliance, and organizational control.

Standard Operating Procedure (SOP): A documented set of step-by-step instructions that describe how a routine process or task should be performed. SOPs ensure consistency, quality, and compliance by standardizing workflows across teams and systems. They are commonly used in regulated industries and operational environments to guide execution and maintain accountability.

Structured Data: Data that resides in fixed fields within a record or file. Relational databases and spreadsheets are examples of structured data.

Unstructured Data: Data that does not reside in fixed locations. Free form text in a word processing document is a typical example.

Workflow Automation: A subset of automation that focuses on coordinating and executing a series of tasks or processes across systems, applications, or teams. Workflow automation maps out the steps in a business or technical process and uses automation to ensure they are carried out in the correct sequence with minimal manual input.

Zero Trust: A security framework that assumes no user, device, or system should be trusted by default, whether inside or outside an organization's network. In computing, zero trust requires continuous verification of identity, strict access controls, and monitoring of all activities to minimize risk and protect sensitive data. (See also: Cybersecurity, IAM, Least-Privilege Access, Privacy).

Works Cited

"5 mindshifts to supercharge business growth." *IBM*, Global C-suite Series, 2025. <https://www.ibm.com/thought-leadership/institute-business-value/en-us/c-suite-study/ceo>. Accessed March 3, 2026.

"Agentic AI." *OpenText*, www.opentext.com/ca/what-is/agentic-ai. Accessed February 10, 2026.

"AI principles." *OECD*, 2024. <https://www.oecd.org/en/topics/sub-issues/ai-principles.html>. Accessed March 6, 2026.

Apotheker, Jessica, Beauchene, Vinciane, de Bellefonds, Nicolas et al. "The Widening AI Value Gap." *Boston Consulting Group*, September 2025. <https://media-publications.bcg.com/The-Widening-AI-Value-Gap-October-2025.pdf>. Accessed March 3, 2026.

"Artificial Intelligence Risk Management Framework (AI RMF 1.0)." *National Institute of Standards and Technology – U.S. Department of Commerce*, January 2023. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>. Accessed March 9, 2026.

"Auckland Transport expands video analytics and realizes its vision of safer roads and more efficient public transportation." *OpenText (Customer Story): Auckland Transport*. <https://www.opentext.com/customers/auckland-transport>. Accessed March 13, 2026.

Awad, Neveen, Serry, Mahmood, and Joe Vasquez. "How Agentic AI Is Transforming Enterprise Platforms." *Boston Consulting Group*, October 13, 2025. <https://www.bcg.com/publications/2025/how-agentic-ai-is-transforming-enterprise-platforms>. Accessed March 3, 2026.

"Bank supports better customer service, speeds request fulfillment and reduces paper trails with OpenText™ Service Management Automation X (SMA)." *OpenText (Customer Story): Indian Ocean Bank*. <https://www.opentext.com/customers/indian-ocean-bank>. Accessed March 13, 2026.

Barrenechea, Mark and Tom Jenkins. *Digital Financial Services*. OpenText, 2016. <https://www.opentext.com/assets/documents/en-US/pdf/opentext-ceo-book-digital-financial-services-en.pdf>. Accessed March 16, 2026.

Barrenechea, Mark and Tom Jenkins. *Digital Manufacturing*. OpenText, 2018. <https://www.opentext.com/assets/documents/en-US/pdf/opentext-ceo-book-digital-manufacturing-en.pdf>. Accessed March 16, 2026.

Barrenechea, Mark and Tom Jenkins. *e-Government or Out of Government*. OpenText, 2014. <https://www.opentext.com/assets/documents/en-US/pdf/opentext-e-government-or-out-of-government-en.pdf>. Accessed March 16, 2026.

Barrenechea, Mark, Jenkins, Tom and David Fraser. *The Anticipant Organization*. OpenText, 2022. <https://www.opentext.com/ca/resources/ceo-thought-leadership/the-anticipant-organization>. Accessed March 16, 2026.

Bell, Shannon. "How to Save \$1 Billion Through Information Management Powered by Automation & AI." *OpenText*, 2025. <https://www.opentext.com/en/media/white-paper/how-to-save-1-billion-ceo-wp-en.pdf>. Accessed March 13, 2026.

Bell, Shannon, Fraser, David and Tom Jenkins. *Enterprise Artificial Intelligence: Building Trusted AI in the Sovereign Cloud*. OpenText, 2025. <https://www.opentext.com/media/ebook/enterprise-artificial-intelligence-building-trusted-ai-with-secure-data-ebook-en.pdf>. Accessed March 16, 2026.

Bennett, Nicole. "100+ AI Statistics Shaping Business in 2026." *Vena*, January 15, 2026. <https://www.venasolutions.com/blog/ai-statistics>. Accessed March 12, 2026.

"Bosch leverages OpenText™ Aviator Lab to explore ways in which AI could streamline investigations, reduce costs, and deliver data-driven insights." *OpenText (Customer Story): Bosch*. <https://www.opentext.com/customers/bosch>. Accessed March 13, 2026.

Brohan, Mark. "Gartner: AI agents will command \$15 trillion in B2B purchases by 2028." *Digital Commerce 360*, November 28, 2025. <https://www.digitalcommerce360.com/2025/11/28/gartner-ai-agents-15-trillion-in-b2b-purchases-by-2028/>. GARTNER is a trademark of Gartner, Inc. and/or its affiliates. Accessed March 12, 2026.

Challapally, Aditya, Pease, Chris, Raskar, Ramesh and Pradyumna Chari. "The GenAI Divide: The State of AI in Business 2025." Project NANDA. MIT, July 2025. https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf. Accessed February 10, 2026.

Chaurasia, Jayesh and Sudha Maheshwari. "Predictions 2025: AI and Automation." *Forrester Research*, 2024.

Chui, Michael. "The state of AI in 2025: Agents, innovation, and transformation." *McKinsey & Company*, November 5, 2025. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>. Accessed February 12, 2026.

Coshow, Tom, Gao, Arnold, Pingree, Lawrence et al. "Top Strategic Technology Trends for 2025: Agentic AI." *Gartner*, October 21, 2024. <https://www.gartner.com/en/documents/5850847>. GARTNER is a trademark of Gartner, Inc. and/or its affiliates. Accessed March 12, 2026.

"Could the CAIO Please Stand Up? The Rising Need for a Chief Artificial Intelligence Officer." *IDC*, September 2024. <https://info.idc.com/rs/081-ATC-910/images/IDC-CAIO-Please-Stand-Up-eBook.pdf>. Accessed March 3, 2026.

"County of Los Angeles automates appeals process with OpenText. Most populous county in the US resolves applicant, employee protests quickly, accurately with eAppeals system powered by OpenText." *OpenText (Customer Story): County of Los Angeles, Department of Human Resources*. <https://www.opentext.com/customers/county-of-los-angeles-department-of-human-resources>. Accessed March 13, 2026.

Covello, Jim, Nathan, Allison, Briggs, Joseph et al. "Gen AI: Too much spend, too little benefit?" *Goldman Sachs Global Investment Research*, June 25, 2024.

Curran, Rowan, Joseph, Leslie, Hopkins, Brian and Craig Le Clair. "Agentic AI Is the Next Competitive Frontier." *Forrester blog*, March 11, 2025. www.forrester.com/blogs/agentic-ai-is-the-next-competitive-frontier/. Accessed February 10, 2026.

"Danish pension and life insurance company automates weekly tests and boosts deployment speed with OpenText DevOps solutions." *OpenText (Customer Story): Velliv, Pension & Life insurance*. <https://www.opentext.com/customers/velliv>. Accessed March 13, 2026.

"DGI Data Governance Framework." The Data Governance Institute. <https://datagovernance.com/the-dgi-data-governance-framework/>. Accessed March 15, 2026.

Dunakhe, Harish, Murat, Melih and Eric Samuel. "The Rise of Agentic AI in the GCC, Riding the Next Wave of AI Revolution." *IDC*, October 2025. <https://www.etisalat.ae/content/dam/eand/assets/img/general/e-aws-whitepaper-final.pdf>. Accessed February 11, 2026.

"Enterprise software leader reduced document types by 96% and prepared for AI innovation in HR with OpenText Content Management for SAP SuccessFactors." *OpenText (Case Study): SAP*. <https://www.opentext.com/customers/sap-hr-document-management>. Accessed March 13, 2026.

"Ethics of Artificial Intelligence." *UNESCO*, November 2021. <https://www.unesco.org/en/artificial-intelligence/recommendation-ethics>. Accessed March 4, 2026.

"EU Artificial Intelligence Act." *European Commission, the European Parliament, and the Council of the European Union*, August 1, 2024. <https://artificialintelligenceact.eu/>. Accessed March 4, 2026.

Fernandez, Dan. "The Developer's Guide to Agentic AI: The Five Stages of Agent Lifecycle Management." *DevOps.com*, October 22, 2025. <https://devops.com/the-developers-guide-to-agentic-ai-the-five-stages-of-agent-lifecycle-management/>. Accessed February 12, 2026.

"Forrester Unveils Top 10 Emerging Technologies For 2025; AI Innovation Shifts From Experimentation To Business Imperative." *Forrester press release*, May 6, 2025. <https://investor.forrester.com/news-releases/news-release-details/forrester-unveils-top-10-emerging-technologies-2025-ai/>. Accessed February 11, 2026.

Francis, Jameel. "Why 85% Of Your AI Models May Fail." *Forbes*, November 15, 2024. <https://www.forbes.com/councils/forbestechcouncil/2024/11/15/why-85-of-your-ai-models-may-fail/>. Accessed March 4, 2026.

"Gartner Hype Cycle Identifies Top AI Innovations in 2025." *Gartner*, August 5, 2025. <https://www.gartner.com/en/newsroom/press-releases/2025-08-05-gartner-hype-cycle-identifies-top-ai-innovations-in-2025>. GARTNER is a trademark of Gartner, Inc. and/or its affiliates. Accessed February 11, 2026.

"Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026, Up from Less Than 5% in 2025." *Gartner*, August 26, 2025. <https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025>. GARTNER is a trademark of Gartner, Inc. and/or its affiliates. Accessed February 11, 2026.

"Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027." *Gartner*, June 25, 2025. www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027. GARTNER is a trademark of Gartner, Inc. and/or its affiliates. Accessed February 9, 2026.

"Gartner Survey Finds Just 15% of IT Application Leaders Are Considering, Piloting, or Deploying Fully Autonomous AI Agents." *Gartner*, September 30, 2025. <https://www.gartner.com/en/newsroom/press-releases/2025-09-30-gartner-survey-finds-just-15-percent-of-it-application-leaders-are-considering-piloting-or-deploying-fully-autonomous-ai-agents>. GARTNER is a trademark of Gartner, Inc. and/or its affiliates. Accessed March 3, 2026.

"Gartner Unveils Top Predictions for IT Organizations and Users in 2026 and Beyond." *Gartner*, October 21, 2025. www.gartner.com/en/newsroom/press-releases/2025-10-21-gartner-unveils-top-predictions-for-it-organizations-and-users-in-2026-and-beyond. GARTNER is a trademark of Gartner, Inc. and/or its affiliates. Accessed February 11, 2026.

Ghosh, Shaona, Simkin, Barnaby, Shiarlis, Kyriacos et al. "A Safety and Security Framework for Real-World Agentic Systems." *NVIDIA*, November 27, 2025. <https://arxiv.org/pdf/2511.21990>. Accessed February 11, 2026.

"Global health technology leader improves maintenance of lifesaving devices and reduces downtime with OpenText™ Analytics Database." *OpenText (Customer Story): Philips Healthcare*. <https://www.opentext.com/customers/philips-healthcare-3>. Accessed March 13, 2026.

Goetz, Michele and Michael O'Grady. "AI Governance Software Spend Will See 30% CAGR from 2024 to 2030." *Forrester blog*, November 13, 2024. www.forrester.com/blogs/ai-governance-software-spend-will-see-30-cagr-from-2024-to-2030/. Accessed February 10, 2026.

Gohring, Nancy. "From risk to reward: The dual reality of agentic AI in the enterprise." *IDC*, December 3, 2025. <https://www.idc.com/resource-center/blog/from-risk-to-reward-the-dual-reality-of-agentic-ai-in-the-enterprise/>. Accessed February 12, 2026.

"Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions." *OECD*, September 18, 2025. https://www.oecd.org/en/publications/governing-with-artificial-intelligence_795de142-en/full-report/ai-in-public-service-design-and-delivery_09704c1a.html. Accessed March 5, 2026.

"How To Lead in the Age of AI Agents." *Salesforce*, July 17, 2025. <https://www.salesforce.com/ap/blog/ai-workforce-management/>. Accessed March 4, 2026.

Horvat, Angel. "Why 80% of Enterprise AI Projects Fail." *AI Readiness*, February 20, 2026. <https://www.aireadi.io/blog/why-80-percent-of-enterprise-ai-projects-fail>. Accessed March 13, 2026.

Higgins, Sam. "From Prompts To Plans: Overcoming The Complexity Gap Between GenAI And AI Agents." *Forrester blog*, July 30, 2025. <https://www.forrester.com/blogs/from-prompts-to-plans-overcoming-the-complexity-gap-between-gen-ai-and-ai-agents/>. Accessed March 6, 2026.

"HR Leaders to Redeploy a Quarter of Their Workforce as Agentic AI Adoption Expected to Grow 327% by 2027." *Salesforce*, May 5, 2025. <https://www.salesforce.com/news/stories/agentic-ai-impact-on-workforce-research/>. Accessed March 4, 2026.

"IBM Study: CEOs Double Down on AI While Navigating Enterprise Hurdles." *IBM Newsroom*, May 6, 2025. <https://newsroom.ibm.com/2025-05-06-ibm-study-ceos-double-down-on-ai-while-navigating-enterprise-hurdles>. Accessed March 13, 2026.

Jarrah, M. H. "Artificial Intelligence and the Future of Work: Human-AI Symbiosis in Organizational Decision Making." *Business Horizons*, 2018, 61(4), 577–586.

Jenkins, Tom. *Managing Content in the Cloud: Enterprise Content Management 2.0*. OpenText, 2011.

Joseph, Leslie. "Announcing Our Evaluation Of The Agent Control Plane Market." *Forrester blog*, December 4, 2025. <https://www.forrester.com/blogs/announcing-our-evaluation-of-the-agent-control-plane-market/>. Accessed February 12, 2026.

Jyoti, Ritu, Lava, Shari, Murat, Melih et al. "IDC FutureScape: Worldwide Artificial Intelligence and Automation 2024 Predictions." *IDC*, October 2023. https://www.idc.com/wp-content/uploads/2025/03/IDC_FutureScape_Worldwide_Artificial_Intelligence_and_Automation_2024_Predictions_-_2023_Oct.pdf. Accessed February 12, 2026.

"Keeping the wheels of justice turning in a global crisis—SMAX enables faster issue resolution and 40% cost reduction." *OpenText* (Customer Story): Court of Justice of the Federal District and Territories (TJDFT). <https://www.opentext.com/customers/court-of-justice-of-the-federal-district-and-territories>. Accessed March 13, 2026.

Lee, Rex. "The Agentic Archetype." *IDC Blog*, December 2, 2025. <https://www.idc.com/resource-center/blog/the-agentic-archetype/>. Accessed February 10, 2026.

Loomis, Amy. "Capitalizing on agentic workflows to enable new work models." *IDC*, September 2025. https://www.idc.com/wp-content/uploads/2025/09/DIR2025_TECHA_AgenticWorkflows_AL.pdf. Accessed February 10, 2026.

McCartney, Ava. "Don't Rush to Appoint a Chief AI Officer." *Gartner*, March 22, 2024. <https://www.gartner.com/en/articles/don-t-rush-to-appoint-a-chief-ai-officer>. GARTNER is a trademark of Gartner, Inc. and/or its affiliates. Accessed March 3, 2026.

"Michelin achieves flexibility and scalability to support their global operation." *OpenText* (Customer Story): Michelin. <https://www.opentext.com/customers/michelin>. Accessed March 13, 2026.

Miller, Michael J. "Riding the AI Whirlwind: Gartner's Top Strategic Predictions for 2025." *PC Magazine*, October 25, 2024. <https://www.pcmag.com/news/riding-the-ai-whirlwind-gartners-top-strategic-predictions-for-2025>. GARTNER is a trademark of Gartner, Inc. and/or its affiliates. Accessed February 22, 2026.

Mueller, Curt, Piasecki, Darryl, El Hoyek, Marie et al. "How COOs maximize operational impact from gen AI and agentic AI." *McKinsey & Company*, March 20, 2025. <https://www.mckinsey.com/capabilities/operations/our-insights/how-coos-maximize-operational-impact-from-gen-ai-and-agentic-ai>. Accessed March 3, 2026.

"OpenText banks on OpenText: Information Management technology leader transforms project management with OpenText™ Project and Portfolio Management for CSO AI team." *OpenText* (Customer Story). <https://www.opentext.com/customers/opentext-ppm-for-cso-ai>. Accessed March 13, 2026.

"OpenText delivers faster insights into consumer sales trends, resulting in 28 percent revenue increases for Sky IT Group clients." *OpenText* (Customer Story): Sky IT Group. <https://www.opentext.com/customers/sky-it-group>. Accessed March 13, 2026.

"OpenText Functional Testing AI capabilities improve regression testing times by 90% and enhance test coverage while aligning with corporate DevOps delivery." *OpenText* (Customer Story): Roche Diagnostics Shanghai. <https://www.opentext.com/customers/roche-diagnostics>. Accessed March 13, 2026.

"OpenText trusts OpenText AI: World leader in information management amplifies employee productivity and enhances customer experiences with AI content management." *OpenText* (Customer Story). <https://www.opentext.com/customers/opentext-ollie-ai>. Accessed March 13, 2026.

Perez, Carlos E. "Why Agentic AI Requires a Radically New Lifecycle." *LinkedIn*, November 7, 2025. <https://www.linkedin.com/pulse/why-agentic-ai-requires-radically-new-lifecycle-carlos-e-perez-qx7xe/>. Accessed February 12, 2026.

"Precision engineering company enables intelligent automation for key back-office processes with OpenText Vendor Invoice Management for SAP Solutions." *OpenText* (Customer Story): KRAMSKI. <https://www.opentext.com/customers/kramski>. Accessed March 13, 2026.

Punjabi, Divesh. "The Agentic Development Life Cycle: How to Manage AI Agents at Scale." *Stack AI*, Feb 6, 2026. <https://www.stack-ai.com/blog/the-agentic-development-life-cycle-how-to-manage-ai-agents-at-scale>. Accessed February 12, 2026.

Raisch, Sebastian and Sebastian Krakowski. "Artificial Intelligence and Management: The Automation–Augmentation Paradox." *Academy of Management Review*, 2021, 46(1), 192–210. https://www.researchgate.net/profile/Sebastian-Raisch/publication/348483816_Artificial_Intelligence_and_Management_The_Automation-Augmentation_Paradox/links/6755c5e1ad10b614ef38f8bd/Artificial-Intelligence-and-Management-The-Automation-Augmentation-Paradox.pdf. Accessed February 10, 2026.

Raji, Inioluwa Deborah, Smart, Andrew, White, Rebecca. N. et al. "Closing the AI accountability gap: defining an end-to-end framework for internal algorithmic auditing." *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, January 27, 2020. <https://dl.acm.org/doi/pdf/10.1145/3351095.3372873>. Accessed February 27, 2026.

Ransbotham, S., Kiron, D., Khodabandeh, S. et al. "The Emerging Agentic Enterprise: How Leaders Must Navigate a New Age of AI." *MIT Sloan Management Review and Boston Consulting Group*, November 2025.

Schadler, Ted, Nagle, Benjamin with Mark Moccia, Callie Smith, and Ian McPherson. "Data Overview: The CIO's Role In AI Success." *Forrester*, February 4, 2026.

Sharma, Ravikant. "Preparing Government for Agentic AI: Data, Governance, and Operating Model for Responsible Adoption." *IDC*, January 26, 2026. <https://www.idc.com/resource-center/blog/preparing-government-for-agentic-ai/>. Accessed February 12, 2026.

Singla, Alex, Sukharevsky, Alexander, Hall, Bryce et al. "The state of AI in 2025: Agents, innovation, and transformation." *McKinsey & Company*, November 5, 2025. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>. Accessed February 12, 2026.

Singla, Alex, Sukharevsky, Alexander, Yee, Lareina et al. "The state of AI in early 2024: Gen AI adoption spikes and starts to generate value." *McKinsey & Company Survey*, May 30, 2024. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-2024>. Accessed February 3, 2026.

"SMAX unifies disparate organizations, delivering service management excellence, cost savings, and great user engagement." *OpenText* (Customer Story): SOCAR Turkey. <https://www.opentext.com/customers/socar-turkey>. Accessed March 13, 2026.

Sukharevsky, Alexander, Krivkovich, Alexis, Gast, Arne et al. "The agentic organization: Contours of the next paradigm for the AI era." *McKinsey & Company*, September 26, 2025. <https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/the-agentic-organization-contours-of-the-next-paradigm-for-the-ai-era>. Accessed February 12, 2026.

Sukharevsky, Alexander, Kerr, Dave, Hjartar, Klemens et al. "Seizing the agentic AI advantage." *McKinsey & Company*, June 13, 2025.

Sunkara, V. L. "KPIs for AI agents and Generative AI: A Rigorous Framework for Evaluation and Accountability." *International Journal of Scientific Research and Modern Technology*, 2024, 3(4), 22–29. <https://doi.org/10.38124/ijsrmt.v3i4.572>. Accessed February 25, 2026.

Taylor, Josh. "Facebook outage: what went wrong and why did it take so long to fix after social platform went down?" *The Guardian*, October 5, 2021. <https://www.theguardian.com/technology/2021/oct/05/facebook-outage-what-went-wrong-and-why-did-it-take-so-long-to-fix>. Accessed March 10, 2026.

"The AI Divide: Why Most AI Initiatives Fail and How the Top 5% Succeed." *Marlabs*, February 17, 2026.

The DAMA® Data Management Body of Knowledge (DAMA-DMBOK®). *DAMA International*. <https://dama.org/learning-resources/dama-data-management-body-of-knowledge-dmbok/>. Accessed March 8, 2026.

"The Top 10 Emerging Technologies Of 2025." *Forrester*, 2025. <https://go.forrester.com/wp-content/uploads/2025/05/Forrester-The-Top-10-Emerging-Technologies-Of-2025-1.pdf>. Accessed February 11, 2026.

Trần, Minh Anh. "An Introduction to Agentic Workflows You Need To Know." *Lollypop UX UI Design Studio blog*, August 8, 2025. <https://lollypop.design/blog/2025/august/agentic-workflows/>. Accessed March 13, 2026.

Valente, Alla and Cody Scott. "The New Chief Artificial Intelligence Officer Role Balances AI Champion And Risk Manager." *Forrester blog*, April 3, 2024. <https://www.forrester.com/blogs/the-new-chief-artificial-intelligence-officer-caio-role-balances-ai-champion-and-risk-manager/>. Accessed February 16, 2026.

Yee, Lareina, Chui, Michael, Roberts, Roger and Stephen Xu. "One year of agentic AI: Six lessons from the people doing the work." *McKinsey & Company*, September 12, 2025. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/one-year-of-agentic-ai-six-lessons-from-the-people-doing-the-work>. Accessed February 10, 2026.

Zhang, Emily. "Why 95% of GenAI projects fail—and why the 5% that survive matter." *Trullion Blog*, September 8, 2025. <https://trullion.com/blog/why-95-of-ai-projects-fail-and-why-the-5-that-survive-matter/>. Accessed February 20, 2023.

Index

A

- Adoption Roadmap, 13, 221, 237–239, 255
- Agency Orchestrator, 68
- Agent Architects, 242, 280
- Agent Control Plane, 208, 217–218, 277, 280, 282, 294
- Agent-Driven Operations, 196–197
- Agentic AI Genome, 1–304
- Agentic AI Genome Map, 145, 280
- Agentic Artificial Intelligence, Agentic AI, 1–304
- Agentic Benefits Eligibility, 122, 165
- Agentic Development Lifecycle (ADLC), 3, 207–211, 215–218, 230, 235
- Agentic Drug Discovery, 121, 146–147
- Agentic Enterprise, 2–295
- Agentic Food Safety, 121, 127
- Agentic Freight Documentation, 121, 134
- Agentic Genome Map, 3, 7, 12, 28, 73, 75, 77, 79, 81, 86–87, 92–93, 95, 119, 133, 138, 150, 155, 228–229, 234
- Agentic Operating Model, 2, 9, 28, 135, 187, 273
- Agentic Patient Triage, 121, 142–143
- Agentic Permitting and Zoning, 122, 174–175
- Agentic Process Control, 122, 151
- Agentic Workflow, 22, 28, 31, 58, 66, 95, 101, 107, 118, 120, 127, 132, 137, 141, 146, 150–151, 161, 165, 168, 171, 173, 175, 179, 188, 241, 247, 284, 289, 294, 296
- AI Agent, 17, 21–22, 34, 40, 46, 53, 58, 70, 73, 80, 109, 112, 118, 156, 175, 183, 188–189, 197, 227, 232–233, 241, 275–276, 278, 280, 283, 292–296
- AI Governance, 20, 85, 189, 204–205, 215, 230, 242, 249, 276, 280, 282, 293
- AI Governance Manager, 242, 280
- AI Ops Engineer, 242, 280
- AI Pilot, 2, 9, 25, 225
- AI Risk Management Framework (AI RMF), 86, 123, 261, 278, 291
- AI-Driven Analytics, 113, 280
- AI-Powered Assistant, 225, 281
- AI-Ready Platform/Stack, 281
- Alabama Gas (Alagasco), 63–65
- Analytics, 7, 16, 42, 52, 84, 87, 97, 104–107, 112–113, 117, 121, 138, 140–142, 166, 168, 176–177, 199, 205, 262, 266, 280–282, 285, 288, 291, 293
- Anomaly Detection, 99, 123, 142, 148, 157, 165, 176, 202
- Anti-Money Laundering, 121, 123
- Application Programming Interface (API), 18, 33, 38, 40, 43–44, 52, 62, 84, 123–124, 143, 165, 170, 198, 208, 210–211, 215, 219, 241, 250, 281, 285–286
- Application Security as a Service (ASaaS), 110, 281
- ArcelorMittal, 35–37
- Archive Center, 117
- Artificial Intelligence (AI), 1–304
- Asset Maintenance, 121–122, 136–137, 141–142, 152, 159–160
- Auckland Transport, 176–177, 291
- Audit, 4–281
- Audit Trails, 26, 38, 44, 47, 59, 78, 84, 95, 145, 149–150, 190–191, 194, 199, 210, 215, 281
- Auditable Access, 281
- Automation, 3–295
- Automotive, 12, 27, 36, 91, 100, 121, 136–137
- Autonomous Asset Integrity, 122, 156
- Autonomous Batch Record Management, 122, 150
- Autonomous Border Risk Assessment, 122, 164–165
- Autonomous Grid Optimization, 122, 159
- Autonomous Predictive Fleet Maintenance, 121, 133
- Autonomous Regulatory Affairs, 121, 145–146

Autonomous Social Housing, 122, 170

Autonomous Treasury Management, 121, 124

B

Banking, 12, 49, 78, 121, 123, 126, 178, 208, 257

Benefit Eligibility, 122, 170

Bosch, 18–19, 291

Bots, 281, 283, 288

Buncombe County, 31–32

Business Intelligence (BI), 281

Business Network, 2, 9, 33, 35–36, 85, 118, 150, 154, 184, 186

Business Process Management (BPM), 60, 167, 281

C

California Consumer Privacy Act, 169

Canopus, 59–61

Care Performance Analytics, 121, 140–141

Case Management, 6, 63, 81, 83, 87, 117, 122–123, 165

California Consumer Privacy Act (CCPA), 169, 282

Centre of Excellence (CoE), 244, 249

Change Management, 26, 29, 115, 123, 212, 229, 237, 242–243, 247–248, 266, 271, 282

Chatbots, 7, 188, 225, 228, 236, 281, 284, 287

Chief AI Officer (CAIO), 29, 44, 55, 65, 73, 85, 93, 119, 179, 193, 206, 221, 230, 235, 237, 242, 251, 263, 273, 282, 284, 292–296

Chief Data Officer (CDO), 4, 29, 44, 55, 65, 73, 77, 85, 93, 119, 179, 184, 193, 206, 221, 230, 237, 251, 263, 273

Chief Financial Officer (CFO), 4, 29, 44, 55, 73, 86, 93, 119, 125, 179, 193, 206, 221, 224–225, 233, 237, 240, 251, 263, 273

Chief Human Resources Officer (CHRO), 4, 29, 44, 55, 73, 93, 119, 179, 193, 206, 221, 230, 237, 240, 251, 263, 273, 282

Chief Information Officer (CIO), 4, 10, 29, 44, 55, 65, 71, 73, 77, 85, 93, 103, 109, 119, 135, 179, 184, 188, 193, 196, 206, 221, 226, 230, 237, 240, 242, 244, 251, 260, 263, 273, 275, 277–278, 282, 295

Chief Operating Officer (COO), 4, 29, 44, 55, 73, 93, 119, 179, 193, 206, 221, 237, 251, 263, 273, 277, 282, 294

Clinical Trial Optimization, 121, 146–147

Cloud, 2, 8, 10, 13, 16, 24, 31, 39, 41, 44, 48, 50–54, 62, 87, 99–100, 105–106, 109–110, 112, 129–131, 136, 152, 158, 160, 168, 182, 184, 191–192, 194–195, 197–204, 206, 231, 270, 272, 281–283, 285–288, 291, 294

Cognitive Computing Era, 54

Compliance, 4–289

Computer Vision, 98, 128, 136, 160, 169, 175, 282–283

Configuration, 24, 48–49, 54, 118, 197, 214, 220, 260, 281–282

Configure, Price, Quote (CPQ), 105, 282

Content Lifecycle Management (CLM), 282

Content Server, 117

Content Substrate, 117, 184, 186, 192, 198–199, 235, 239, 248, 251

Contextual Intelligence, 247, 282

Contract Management, 122, 131, 166–167

Control Plane, 21, 30, 51, 54–55, 76, 131, 143, 181, 184, 189, 191, 198, 200–201, 208, 217–218, 230, 243–244, 249, 253, 259, 267, 277, 280, 282, 294

Copilot, 8, 34, 109, 182–183, 225, 227–228, 232, 236, 281

County of Los Angeles, Department of Human Resources, 103, 292

Court Case Management, 63

Court of Justice of the Federal District and Territories (TJDFT), 167–168, 294

Credit Underwriting, 121, 124

Customer Acquisition Cost (CAC), 104, 282

Customer Experience (CX), 78, 123–124, 136, 295

Customer Relationship Management (CRM), 9, 18, 20, 22, 33, 37, 39, 84, 96, 104–105, 112, 282, 284

Customs Clearance, 121, 134

Cybersecurity, 42, 87, 164, 226, 267, 282–283, 287, 290

D

DAMA® Data Management Body of Knowledge (DAMA-DMBOK®), 271, 279, 296

Data Exfiltration, 42

Data Governance, 21, 25, 43–44, 46, 55, 96, 189, 224, 270–271, 279, 281, 292

Data Governance Institute, 271, 279, 292

Data Privacy, 26, 170, 267, 282–284, 288

Data Quality, 29, 31, 46, 73, 154, 193, 208, 224, 229, 235, 237, 255, 262

Data Residency, 140, 191, 194, 200, 206, 262, 272, 283, 289

Data Security, 282–283

Data Silos, 25, 224, 283

Data Sovereignty, 51, 130, 164, 191, 243, 280, 283, 289

Davigel, 49–50

Decision Agent, 124, 170

Deep Learning, 111, 159, 283, 286

Defensibility, 3–4, 13, 38, 44, 46, 49, 124, 187, 195, 200–201, 204, 225–226, 228, 239, 249, 251–252, 257–258, 260–263, 269

Deployment, 3–292

DevOps, 276–277, 292, 295

Digital Twin, 92, 127, 142, 151, 159, 283, 285

Digital Workforce, 3, 18–19, 59, 68, 77, 80–81, 283

Directive on Automated Decision-Making, 164

Discovery, 6, 18, 27–28, 109, 115, 117, 121, 146–147

Disruption Management, 121, 132–133

Dynamic Route Optimization, 121, 132–133

E

e-Discovery, 27

Edge Computing, 157, 283

Electronic Data Interchange (EDI), 36, 100, 134, 154, 283

Electronic Health Record (EHR), 140–141, 143, 146

Encryption, 51, 283

Energy and Utilities, 13, 122, 159

Energy Optimization, 122, 151

Enterprise Architecture, 4, 55, 93, 184, 199, 227, 273, 289

Enterprise Artificial Intelligence (EAI), 2, 10, 16, 25, 31, 53, 143, 182, 191, 283

Enterprise Content Management (ECM), 52, 60, 95, 184, 186, 284, 294

Enterprise Information Management (EIM), 117–119, 148–149, 184, 186, 204, 284

Enterprise Orchestrator, 62, 68, 83, 89, 164

Enterprise Resource Planning (ERP), 9, 18, 20, 22, 33, 39, 64, 67, 84, 88, 91, 96, 99, 105, 111–112, 131, 136, 146, 163, 185–187, 197–198, 203, 284

Enterprise Service Management (ESM), 118, 158, 168, 284

Environmental Reporting, 122, 161

Environmental, Social, and Governance (ESG), 62, 97, 161, 185

European Court of Human Rights (ECHR), 81–83, 170

European Union, 140, 145, 234, 261, 275, 278, 284, 292

European Union AI Act (EU AI Act), 86, 109, 123, 164, 234

Executive and Board Agents, 12, 95

Experimentation Loop, 214

F

Facebook, 259, 278, 296

Fallout Rates, 248, 253, 284

Federal Information Security Management Act (FISMA), 164

Feedback Loop, 18, 104–105, 171, 216, 219, 241

Finance, 4, 6, 12, 22, 26–28, 57–58, 62, 65, 76–77, 79, 87, 93, 95–99, 118, 121, 123, 138, 187, 195, 200, 240, 246, 257, 272, 284

Finance and Audit Agents, 12, 95, 97, 138

Financial Crime, 115–116, 121, 123

FinOps, 109, 233

Food and Beverage Manufacturing, 12, 121, 127

Food and Drug Administration (FDA), 128, 147–148

Food Safety Modernization Act (FSMA), 128

Fraud, 6, 17, 87, 98–99, 116, 121, 123, 138, 170, 257–258, 285–286, 288

Fraud Detection, 6, 87, 98, 121, 138, 257–258, 286, 288

G

GenAI Divide, 25–26, 182, 195–196, 204, 224, 275–277, 292

General Data Protection Regulation (GDPR), 67–68, 85, 115, 282, 284

Generative AI (GenAI), 15–17, 19–20, 25–26, 57, 104, 117–118, 166, 182, 188, 195–196, 202, 204, 223–228, 230, 232, 234, 236–237, 266, 275–277, 283–284, 292, 294, 296

Geopolitical, 54, 111–112, 123, 132, 191

Governance, 3–295

H

Health Insurance Portability and Accountability Act (HIPAA), 85, 284

Health, Safety, and Environment (HSE), 122, 157, 285

Healthcare, 6, 12, 26, 49, 53, 86–87, 112–113, 121, 140–143, 169, 178, 200, 293

Hub-and-Spoke Model, 244–245

Human Capital Management (HCM), 67, 284

Human Oversight Owner, 85, 88, 242, 284

Human-in-Command, 30, 44–45, 47, 55, 76, 117, 122, 131, 135, 150, 159, 164, 169, 179, 188, 198, 213–214, 219, 235, 239, 243, 246–247, 250, 253, 267, 280

Human-in-the-Loop, 16, 20, 47, 125, 170, 210, 218, 266–267

Human Resources (HR), 4, 6, 8, 12, 22, 29, 44, 65, 67–68, 76–77, 79–80, 93, 95–96, 101–104, 118–119, 189–190, 193, 195, 230, 240–241, 243, 272, 275, 284, 292, 294

Hybrid Cloud, 202, 285, 288

Hybrid Model, 53

Hyperscaler, 53, 199, 204, 227, 285

I

Identity and Access Management (IAM), 33, 81, 189, 250, 268, 285, 290

In Silico, 146, 285

In-jurisdiction Hosting, 201

Indian Ocean Bank, 125–126, 291

Information Governance, 118

Intake Agent, 124, 165–166, 170, 172, 174

Integration, 10–286

Intelligence Layer, 77, 96–98, 101, 104–105, 107, 111, 114, 124–125, 128–129, 133–134, 137–138, 142–143, 147–148, 151–152, 156–157, 159–161, 165–167, 169–171, 175–176, 266, 285

Intelligent Aftermarket Parts, 121, 137

Intelligent Batch Record Review and Release, 121, 147–148

Intelligent Capture, 117, 131

Intelligent Citizen Service Resolution 311, 122, 173

Intelligent Federal Procurement, 122, 166

Intelligent Production Optimization, 122, 155

Intelligent Service Request, 122, 171, 174

Intelligent Yield Optimization, 121, 128

International Traffic in Arms Regulations (ITAR), 85, 285

Internet of Things (IoT), 111, 127–128, 133, 136, 148, 156–157, 161, 176, 185, 283, 285

Interoperability, 38–39, 41, 166, 170, 175, 192, 249, 286

Invisible Internet Project (I2P), 91, 285

Invoice Processing, 63, 88–90, 187

IT Asset Management (ITAM), 158, 168

IT Service Management (ITSM), 24, 126, 158, 168

IT, Security, and Compliance Agents, 12, 95, 107

J

JavaScript Object Notation (JSON), 88, 285

K

Key Performance Indicators (KPIs), 81, 96, 210–211, 214, 218–219, 239, 242, 248, 253–254, 262, 285, 296

Know Your Customer/Client (KYC), 78, 124, 285

Knowledge Discovery, 117

KRAMSKI, 90–91, 295

Kubernetes, 48, 286

L

Laboratory Information Management System (LIMS), 146, 150–151, 286

Large Language Model (LLM), 16, 18, 20, 25, 85, 104, 114, 147, 166, 286

Least-Privilege Access, 41, 109, 188, 202, 219, 286, 290

Legacy Systems, 25, 33, 43, 102, 164, 169, 173, 199, 286

Legal, Risk, and Records Agents, 12, 95, 114

Liquidity Forecasting, 121, 124

Load Balancing, 122, 159, 286

Loan Origination, 63, 121, 124

Logistics, 4, 6, 12, 46, 50, 111, 121, 130–134, 136–137, 153, 156, 160, 176, 185–187, 196–197, 216–217

Logistics Optimization, 6, 121, 137

M

Machine Learning (ML), 16, 49, 143–144, 148, 168, 266, 280–281, 283–287, 289

Machine-Generated Data, 34

Manufacturing Execution System (MES), 136, 150–151, 286

Mean Time to Detect (MTTD), 259, 287

Mean Time to Repair (MTTR), 108, 259, 287

Metadata, 31, 33–34, 38–40, 43–44, 54, 84, 95, 97, 117, 119, 184, 186, 192–193, 196–197, 201, 229, 235, 271, 286

Michelin, 99–100, 294

Microservices, 33, 108, 287

Middleware Layer, 286

Mission Orchestrator, 68

Model Context Protocol (MCP), 84, 286

Modernization, 24, 128, 162, 197, 254–255, 286

Monolithic, 58, 61, 63, 94, 286

Multi-Cloud, 13, 54, 191, 194, 198–199, 206, 287

Multi-Region, 201, 287

N

Named Entity Recognition (NER), 165, 287

National Institute of Standards and Technology (NIST), 86, 123, 261, 278, 291

Natural Language Processing (NLP), 124, 134, 138, 148, 171, 283, 286–287, 289

Natural Language Understanding (NLU), 101, 143, 157, 168, 174, 281, 287

Nested Orchestration, 56, 59, 61, 69, 71–72, 77, 228–229, 234

NIST AI RMF, 86, 123

O

Oil and Gas, 13, 122, 155, 158

On-Premises (On-Prem), 41, 52, 131, 198–199, 204, 285–286, 288

OpenText, 1, 10–11, 53, 230–232, 240, 243–244, 246–247, 254–255, 259, 291–296

Operations, Supply Chain, and Facilities Agents, 12, 95, 111

Optical Character Recognition (OCR), 63–65, 98, 124, 134, 144, 148, 169–170, 287

Orchestration, 2-287

Orchestration Layer, 2, 4, 17, 28, 76, 105, 121–122, 130, 135, 159, 178, 184, 187, 192, 198, 212, 217, 231, 235, 244, 263, 267, 270, 274, 280, 287

Orchestrator, 8-287

Org Chart, 12, 56, 58–59, 62, 71–72, 83, 266

Organization for Economic Cooperation and Development (OECD), 271, 278–279, 291, 293

P

Peer-to-Peer (P2P), 91, 119, 287

Permissions, 38, 40, 44, 63, 65, 71, 79, 81, 97, 108–109, 118–119, 199–201, 211–213, 218–219, 268, 285–287, 289

Personal Information Protection and Electronic Documents Act (PIPEDA), 288

Personalization, 104, 285

Personally Identifiable Information (PII), 86, 165–166, 288

Pharmaceutical, 13, 121, 145–147, 149

Philips Healthcare, 112–113, 142, 293

Pilots, 2–3, 5, 9, 13, 20, 25, 28–29, 39, 73, 95, 188, 191, 193, 202, 204, 223–226, 228–230, 232, 236–238, 243, 248, 251

Policy Engines, 189, 201

Predictive Asset Maintenance, 122, 152, 160

Predictive Maintenance, 4–6, 111, 113, 122, 133, 156

Predictive Manufacturing, 121, 136–137

Predictive Medical Imaging Asset Maintenance, 121, 141–142

Predictive Municipal Infrastructure Maintenance, 122, 175–176

Predictive Procurement, 121, 129

Privacy, 9, 26, 54, 87, 96, 115, 118, 140, 164–166, 169–170, 173, 189–190, 211–212, 219, 226, 243, 267, 270, 272, 282–285, 288, 290

Privacy-by-Design, 165, 288

Private Cloud, 62, 199, 203, 285, 288 Process and Case Orchestrator, 63 Process Group, 62, 65, 83, 85, 89

Process Management, 60, 164, 167, 281, 288

Process Manufacturing, 13, 122, 150

Procure-to-Pay, 63–64, 68, 89–90, 97, 185, 229, 246

Production Planning, 121, 129–130

Productivity, 9, 27, 29, 44, 46–47, 56–57, 59, 64, 73, 83, 87, 101, 167, 182–183, 188, 204, 232, 234, 246–247, 253–255, 258, 272, 295

Program Orchestrator, 62, 65

Project Portfolio Management (PPM), 231, 288, 294

Prompt (and Interaction) Engineer, 288

Prompt Engineering, 205

Public Cloud, 199, 204, 288

Public Sector - Federal, 13, 122, 164

Public Sector - Local/Municipal, 13, 122, 173

Public Sector - State/Provincial, 13, 122, 169

Q

Quality Assurance, 121, 136–137, 148–151, 202–203, 216

Quality Control, 121, 128, 145–147, 179

R

Recovery Time Objective (RTO), 259, 288

Recruiting, 69

Reference Architecture, 3, 181, 184, 198, 228, 244, 249

Regulatory Compliance, 5, 67, 110, 115, 122, 124, 127, 135, 140, 150, 155, 157,

161, 203, 226, 262, 281, 286, 289

Repository, 22, 31, 33, 42, 95, 149, 168, 184, 186, 289

Reservoir Management, 122, 155

Retrieval-Augmented Generation (RAG), 16, 34, 43, 102, 105, 182, 184, 212, 289

Return on Investment (ROI), 19, 55, 63, 73, 79, 86, 93, 119, 125, 179, 193, 196, 218, 224–225, 232–234, 236–237, 251–253, 256–259, 262–263

Risk, 3–296

Roadmap, 3, 5, 13, 221, 229–230, 237–239, 248–251, 255

Robotic Process Automation (RPA), 254, 288

Roche Diagnostics Shanghai, 144, 295

Roles, 8–10, 26, 29, 41, 55, 59, 61, 69, 73, 80, 85, 102, 108–109, 179, 205, 228, 237, 239–240, 242, 246, 251, 263, 266, 270, 273, 284

Rules Engine, 289

Runtime, 184, 188–189, 192, 201–202, 212–213, 215, 217, 228, 230

Runtime Optimization Loop, 215, 228

S

Safe Food for Canadians Regulations (SFCR), 128

Sales and Marketing Agents, 12, 95, 104

SAP, 64, 67–68, 292, 295

Scalability, 23, 38–39, 41, 134, 203, 237, 285–286, 288–289, 294

Scaling, 3, 5, 22, 48–49, 68–70, 93, 202, 220, 225, 230, 234, 237, 239, 243, 248–250, 253, 270, 274, 280–281,

286–287, 289

Security, 6–293

Sentiment Analysis, 87, 171, 287, 289

Service Level Agreement (SLA), 107–109, 112, 115, 188, 230, 248, 253, 256, 259, 289

Service Oriented Architecture (SOA), 289

Single Source of Truth, 35, 96, 125, 186, 231–232

Sky IT Group, 105–107, 295

SOCAR Turkey, 157–158, 296

Software Development Lifecycle (SDLC), 207–208, 211–212

Sovereign Cloud, 2, 10, 16, 31, 44, 51, 53, 136, 152, 160, 182, 184, 191–192, 195, 198–200, 270, 272, 291

Sovereign Zone, 53, 62, 165, 170, 174, 199, 289

Sovereignty, 30, 44–45, 50–55, 76, 78, 84, 95, 130, 164–165, 191, 195, 200–201, 235, 243, 253, 270, 272–273, 280, 282–283, 285, 289

Sponsorship, 239–240, 246

Standard Operating Procedure SOP, 33, 149, 289

Structured Data, 88, 131, 170, 177, 283, 285, 290

Submission Management, 121, 145–146

Supervisory Control and Data Acquisition (SCADA), 156, 289

Supply Chain, 4, 6, 9, 12, 20, 26, 33, 35, 65, 67, 76–77, 79, 83, 85, 89, 95, 97, 100, 111–112, 118, 121, 127–130, 132–133, 136, 138, 148, 150, 154, 156, 195, 227, 284

Suspicious Activity Report (SAR), 123, 258

T

The Cloud, 8, 39, 50, 52, 200, 231, 282, 294

Traceability Compliance, 121, 127

Transformation, 3–295

Transportation, 6, 12, 53, 121, 132–134, 155, 169, 177, 291

U

Unstructured Data, 117, 283, 290

V

Value Chain, 62, 68, 78, 86, 121, 141, 145, 179, 269, 274

Velliv, 202–203, 292

Vendor Invoice Management (VIM), 64–65, 91, 295

VW Finance, 27–28

W

Warranty Claims Processing, 121, 138

Workflow Automation, 5, 78, 281, 290

Z

Zero Trust, 184, 290

The Agentic AI Genome

Artificial intelligence is entering a new phase. No longer limited to generating insights or assisting with decisions, intelligent systems are beginning to execute work—retrieving information, coordinating workflows, enforcing policies, and taking action across enterprise systems. This shift marks the emergence of the **agentic enterprise**, where humans define intent and governance while networks of specialized AI agents carry out complex tasks at machine speed.

But scaling this new model requires more than powerful algorithms. It requires an architectural foundation that connects trusted data, orchestrated agents, governance frameworks, and human oversight into a cohesive operating system for the enterprise. It requires the **Agentic Genome Map**.

Drawing on real-world enterprise experience and industry research, *The Agentic AI Genome* presents a **practical blueprint** for designing, delivering, and governing agentic AI at scale. It shows how organizations can move beyond isolated AI pilots to a coordinated operating model—embedding intelligent agents into business workflows while maintaining sovereignty, accountability, and control. From governance frameworks to measurable business value, this book equips leaders to master the transition to intelligent operations.

The nature of business is evolving. It's time to rewrite the DNA of your enterprise.