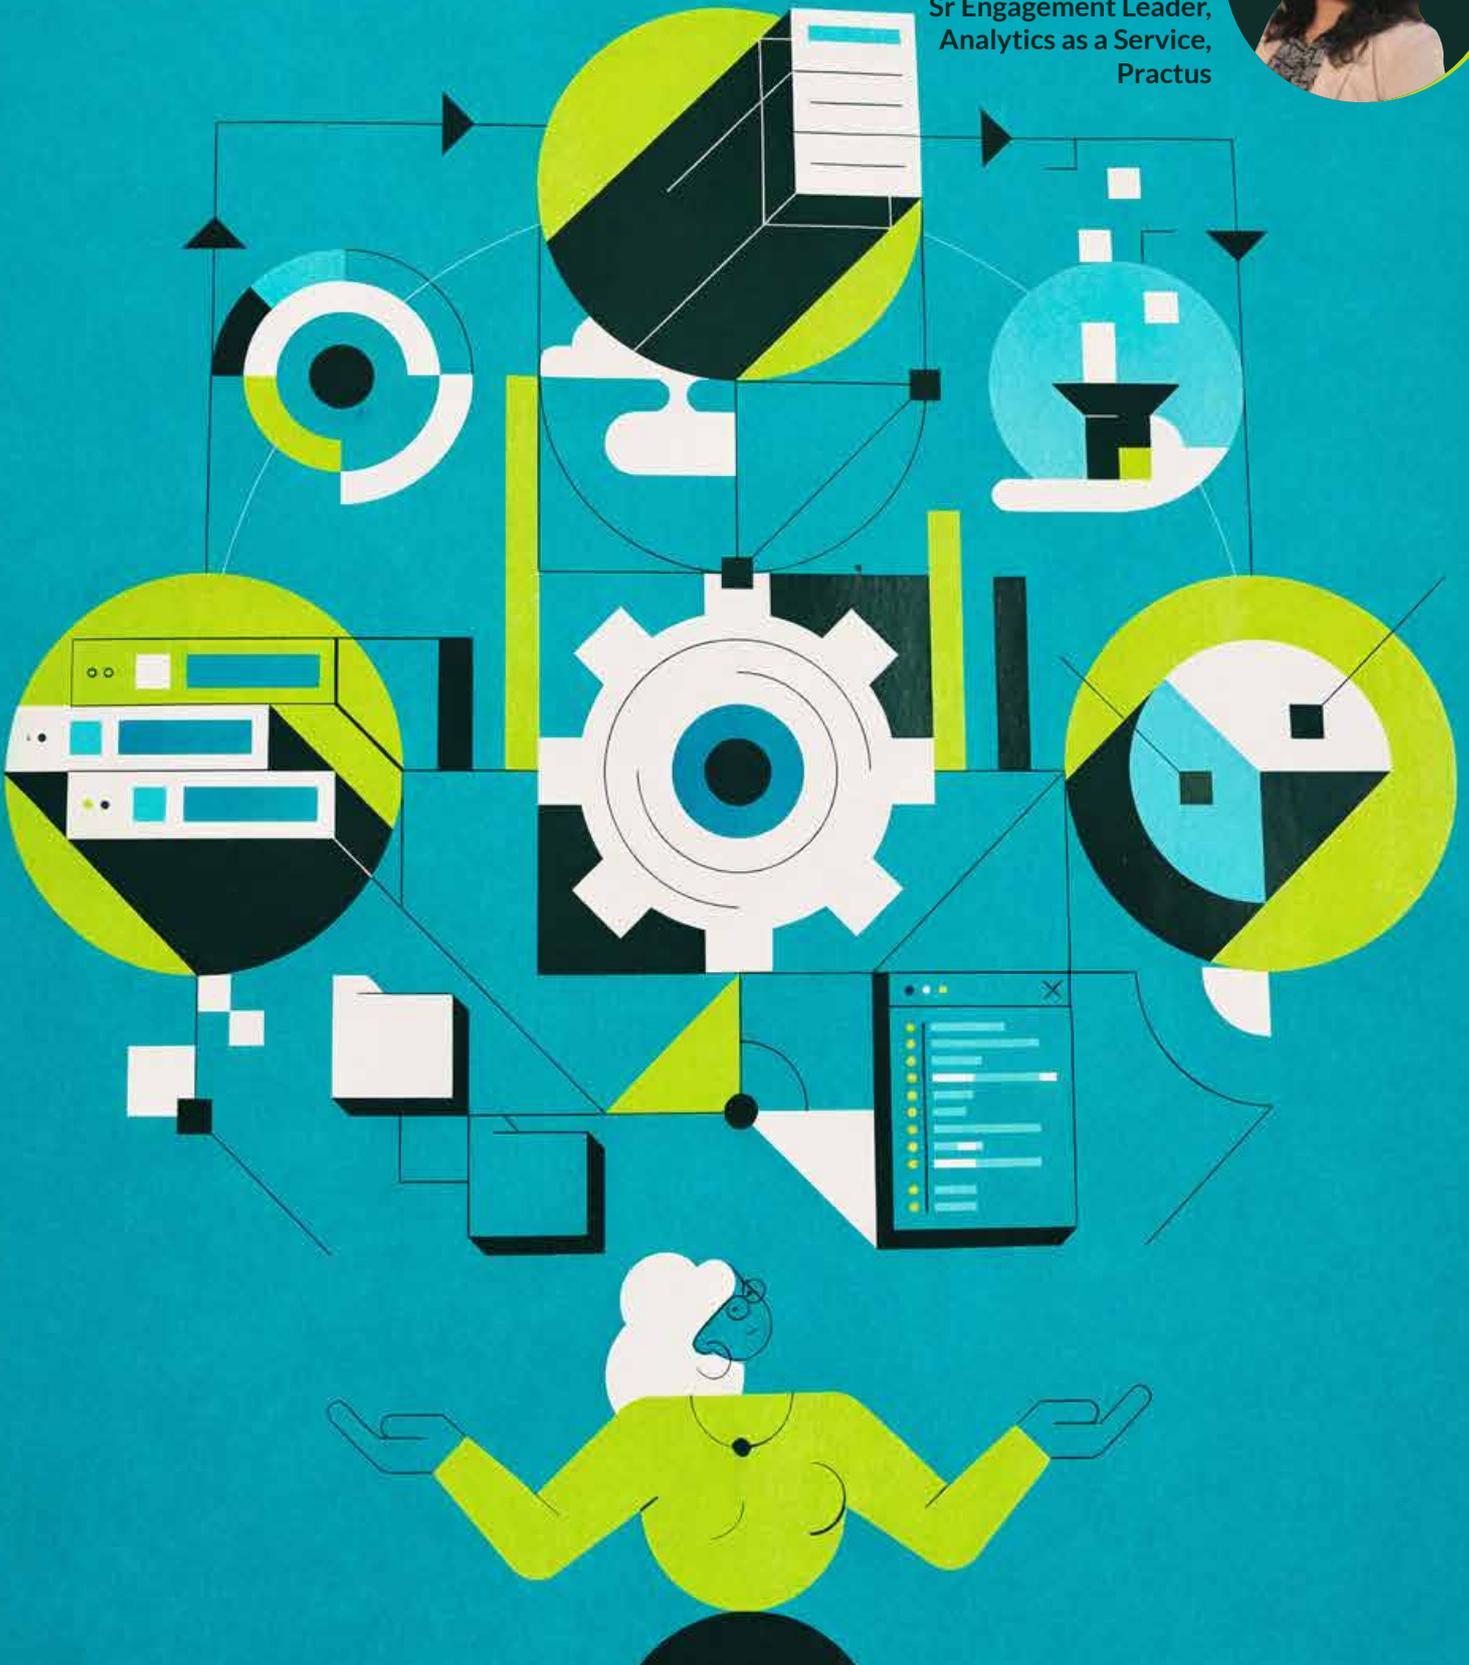


A Maturity Framework for Enterprises Embracing Next-Gen AI

Priyanka Baram
Sr Engagement Leader,
Analytics as a Service,
Practus





Executive Summary

Agentic AI marks a transformation from the intelligent automation of a few processes to the autonomous execution of decisions within enterprise workflows. Readiness for this level of cognitive technology implementation requires disciplined governance, clear decision rights, an integrated architecture, and operating models designed for supervised autonomy.

This paper introduces an Agentic Maturity Matrix as a framework to assess organizational preparedness across leadership intent, data coherence, technical spine, human oversight, and governance design. It outlines how enterprises can transition to AI agent systems by building controlled autonomy.

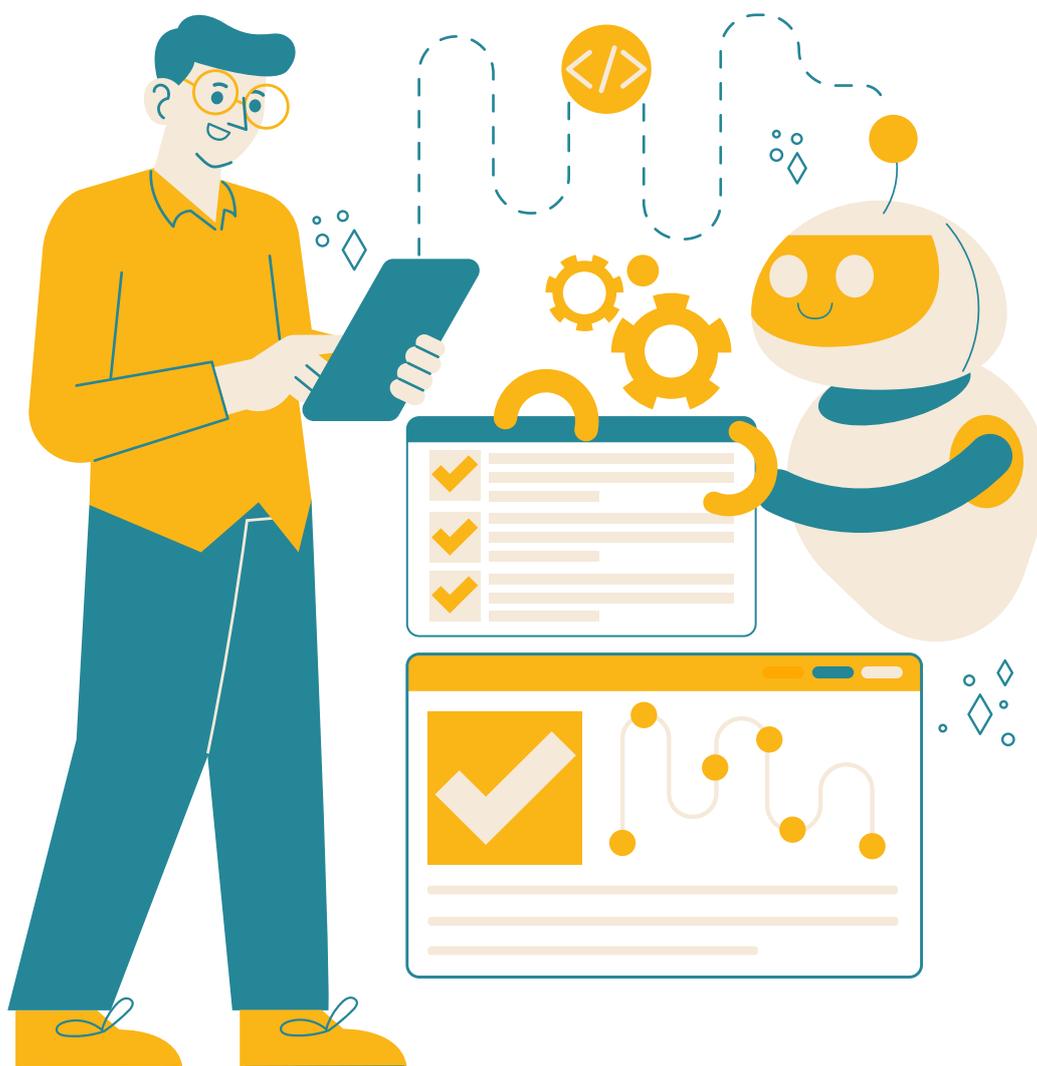
As business leaders across industries adopt next-gen AI, they can assess their organization's preparedness to delegate judgment responsibly and use a defined roadmap to activate autonomous capabilities sustainably.



Introduction

Enterprise AI has entered the active execution phase. Organizations – large, medium, small – are embedding intelligent systems into their operational workflows, customer engagement models, supply networks, and financial controls. They are cautiously implementing agentic tools that can interpret context, initiate actions, and adapt within defined parameters.

Findings from the [Stanford Institute for Human-Centered Artificial Intelligence](#) show sustained growth in enterprise AI adoption, with companies integrating AI into core business functions rather than treating it as a peripheral innovation initiative. Research activity on multi-agent systems and autonomous coordination is also expanding – this signals a maturing field for enterprise-grade usage





The Shift Has Already Begun

For years, enterprises have invested in dashboards, predictive models, and automation. But the next wave of AI is not just about prediction — it is about autonomous decision-making and execution.

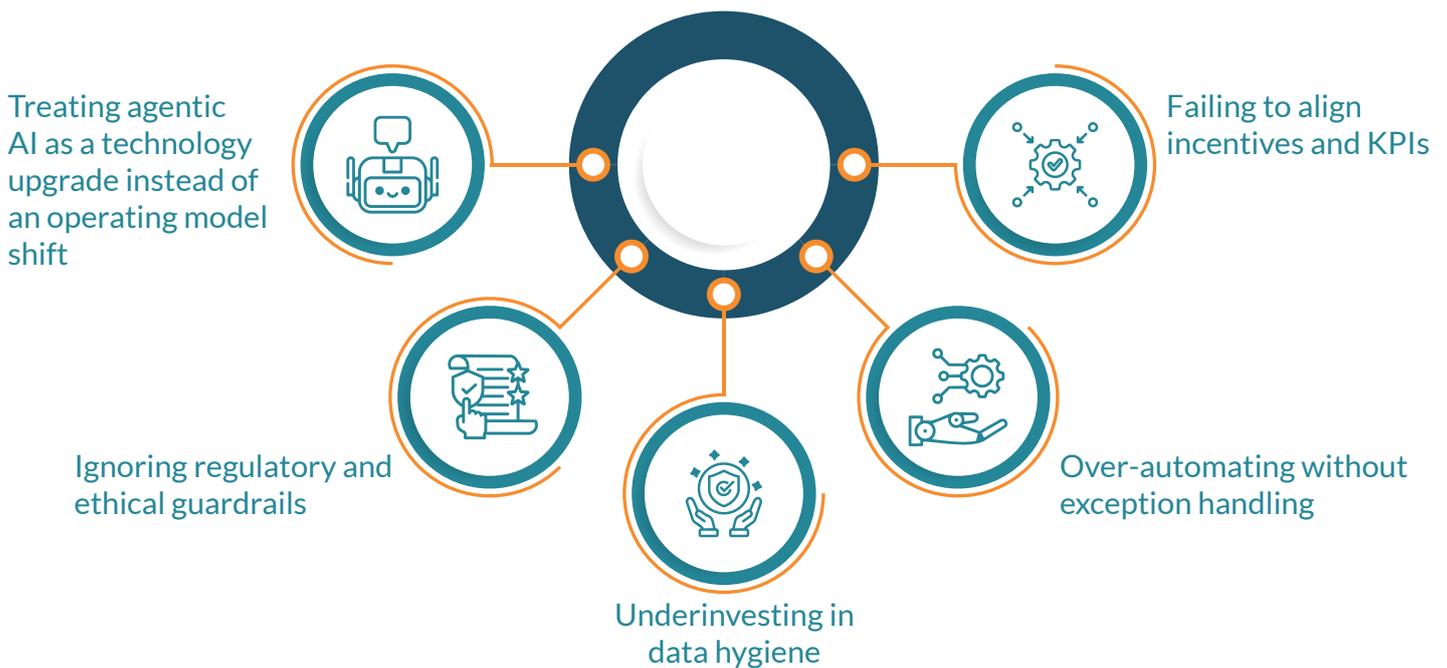
Agentic AI systems do not simply recommend actions; they plan, decide, act, and continuously learn within defined guardrails.

The question for CXOs is no longer **“Should we adopt AI?”**
It is: **“Is our enterprise ready to trust AI to act?”**

Based on our work across retail, pharma, healthcare, and telecom, we see a consistent pattern: organizations rush toward agentic solutions without first building foundational maturity. The result? Pilot fatigue, governance concerns, and unrealized ROI.

While the need for autonomous AI systems looks urgent, capital accountability also matters. CFOs need technology investments to provide tangible ROI via revenue expansion, cost discipline, resilience, and risk mitigation. Instead of merely supporting isolated productivity improvements, agentic systems must also help enterprises achieve gains in responsiveness, adaptability, and growth. Preparing for value-generating AI applications requires assessment of the current state and charting a disciplined path toward scalable, responsible autonomy.

Common Pitfalls We See





The Pillars of Agentic Readiness

Agentic readiness is an organization's capacity to implement autonomous, goal-driven AI systems in a controlled, value-generating manner. It is an integrated capability across five dimensions.



Planned Orientation

As AI initiatives are tied to measurable business outcomes, leadership must define processes that autonomous agents can manage, where human authority is essential, and how success will be quantified.



Data Integrity & Accessibility

Data to feed AI must be clean, accurate, unbiased, current, governed, and interoperable across all systems. When data definitions are standardized enterprise-wide, everyone can operate from the same source of truth. Information is available in real time, allowing agents to act instantly rather than work on after-the-fact analysis.



Technical Architecture

Modular, API-enabled infrastructure is a prerequisite for agentic AI, enabling seamless orchestration across applications. The technical foundation must support monitoring, audit trails, and scalable model deployment without brittle integrations.



Organizational Adaptability

Employees need to be trained to work alongside AI systems rather than treating them as tools. Decision ownership must be clearly defined when agents act independently. Teams collaborate cross-functionally to design, monitor, and refine AI-enabled workflows.



Governance & Risk Controls

Defined escalation thresholds, explainability standards, compliance mapping, ethics, and accountability frameworks explain how decisions are supervised. Supervision must account not only for individual AI models but also for how multiple agents interact, influence one another, and collectively shape business outcomes.

Agentic-readiness is multidimensional. Strength in one area cannot compensate for fragility in another. An organization with advanced architecture but poor governance amplifies risk. Sharp strategy without reliable data constrains execution. Sustainable autonomy needs balanced maturity through all five dimensions.



Governance in Autonomous Systems

Governance must be a core design discipline, not an afterthought. Decision rights need explicit allocation: what agents execute independently, what requires human validation, and where override authority applies. Traceability must be embedded at the system level, capturing the full decision flow across agents, data, and human interventions. Frameworks such as the EU AI Act and NIST AI RMF demand systemic accountability. Exception-handling protocols and AI oversight committees ensure autonomy stays within defined guardrails.

Building the Technical Spine

Agentic AI demands an integrated backbone connecting data, models, agents, and operational systems into a coherent decision fabric. A decision-instrumentation layer supports context persistence, state management, and dynamic task allocation. A shared memory substrate grounds agents in business policies and domain constraints. Architectural observability captures drift, latency, and confidence for continuous calibration. Resilience – through fallback logic and controlled degradation – must be engineered in from the start.

Organizational Redesign for an Agentic Enterprise

Operating models must evolve so human judgment keeps pace with machine-led execution. Leaders shift from directing tasks to defining intent, constraints, and performance thresholds. New roles emerge at the intersection of technology and operations: agent trainers, system supervisors, policy designers, and decision auditors. Metrics expand beyond siloed productivity to capture system contribution and end-to-end outcomes. Enterprises must normalize collaboration with autonomous systems, building trust through transparency and visible accountability.





The Agentic-Readiness Maturity Framework

We define five maturity layers that determine enterprise readiness.

1. Data & Intelligence Foundation

Most organizations are still battling fragmented systems, inconsistent master data, and limited real-time visibility.

Maturity Indicators:

1 Unified data architecture	3 Data governance and lineage tracking
2 Near real-time data pipelines	4 High model explainability

For instance, a leading omnichannel retailer implemented an AI-based replenishment model. Forecast accuracy improved by 12%, but execution failed due to a 24-hour delay in store-level inventory feeds. Agentic systems require operational-grade data, not analytical-grade lagged data.

On similar lines, a telecom operator that was deploying churn prediction using AI realized that customer 360 views were inconsistent across the prepaid and postpaid segments. Without unified data, autonomous retention campaigns risked misfiring.

2. Process Standardization & Digitization

Agentic systems thrive in well-defined, digitized workflows. Chaos cannot be automated.

Maturity Indicators:

1 Standard Operating Procedures (SOPs) documented	3 Clear exception handling paths
2 API-enabled systems	4 Measurable process KPIs

In commercial pharma operations, promotional planning across territories varied widely by region. An agentic AI system designed to optimize field force allocation struggled because its execution logic differed from state to state. Standardization was required before autonomy.

Similarly, hospital revenue cycle management AI could identify billing errors. However, autonomous correction was not possible due to non-standard coding practices across departments.

3. Decision Governance & Risk Controls

Agentic AI must operate within guardrails. Enterprises often underestimate regulatory, ethical, and reputational risk.

Maturity Indicators:

1 Defined decision-rights matrix	3 Audit logs and explainability layers
2 Escalation triggers for high-risk decisions	4 Regulatory compliance alignment

A healthcare provider network tested autonomous patient scheduling optimization. While efficiency improved, fairness concerns arose when high-margin procedures were prioritized.

Governance frameworks were redesigned to ensure ethical balance.



4. Technology & Orchestration Layer

Agentic AI requires seamless orchestration across CRM, ERP, supply chain, and digital platforms.

Maturity Indicators:

1 Microservices-based architecture	3 Cloud-native scalability
2 API-first integration strategy	4 Secure access controls

An AI pricing agent in the retail industry dynamically adjusted online prices based on demand elasticity. However, without integration to offline POS systems, channel conflict emerged. True agentic capability demands cross-channel synchronization.

For Pharma, Supply chain risk agents monitoring API (Active Pharmaceutical Ingredient) shortages must connect directly to procurement systems to automatically trigger alternate sourcing.

5. Organizational & Cultural Readiness

Technology maturity without cultural readiness leads to silent resistance.

Maturity Indicators:

1 AI literacy across leadership	3 Performance metrics aligned with AI augmentation
2 Defined human-in-the-loop roles	4 Change management programs

Across industries, we observe the same pattern: frontline teams fear loss of control.

Successful enterprises reposition agentic AI as a co-pilot, not a replacement.





Agentic Maturity Matrix

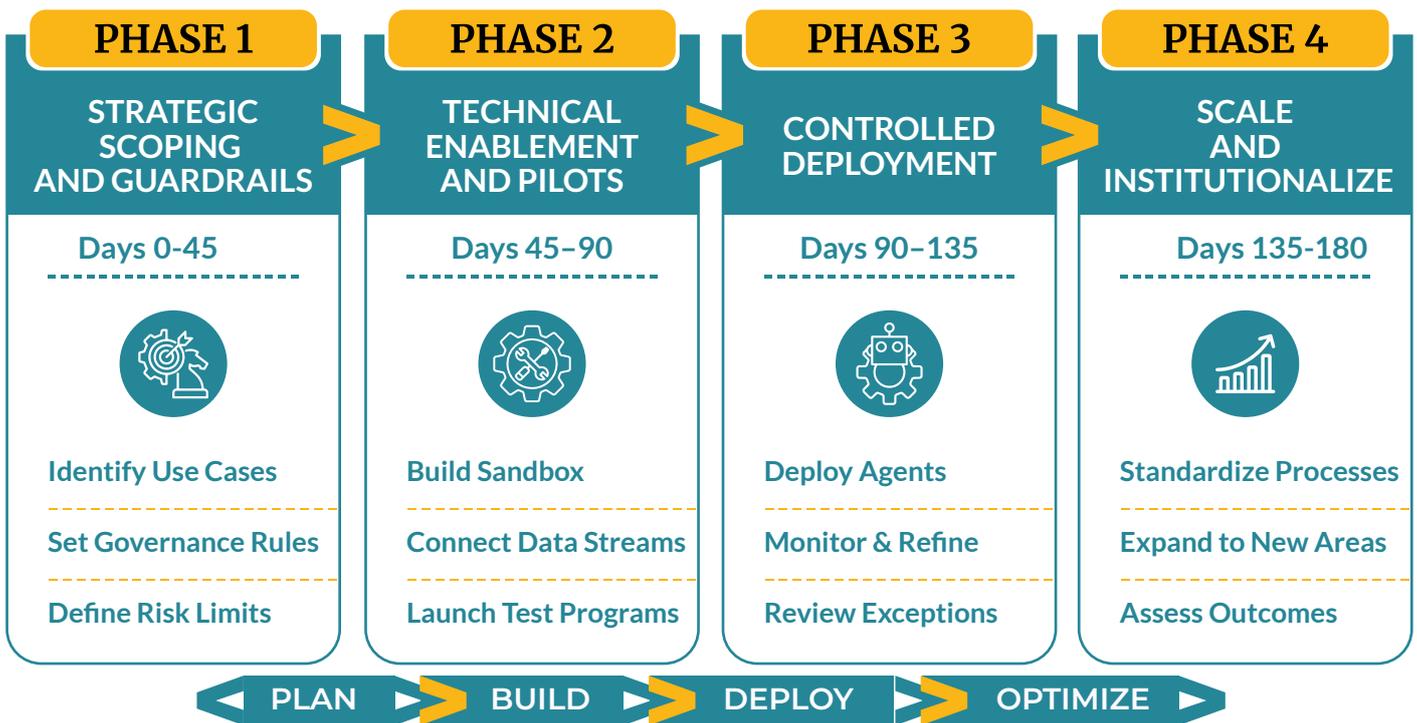
Agentic transformation does not occur in a single leap. Enterprises evolve through identifiable stages of capability development. This Agentic Maturity Matrix provides a structured view of the progression, mapping five levels of maturity spanning core enterprise dimensions.



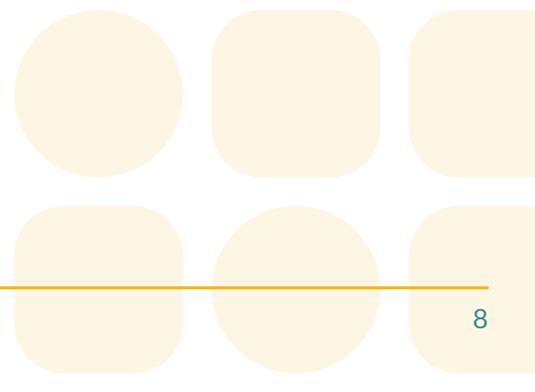
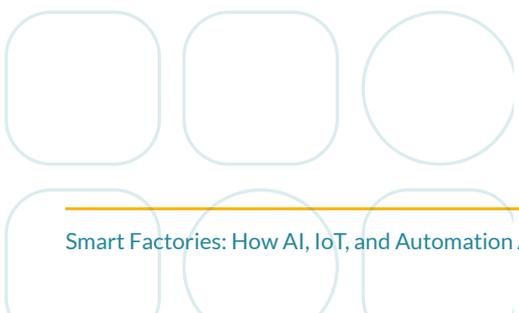
Becoming an agentic enterprise is not a single leap: it is a progression through four distinct stages, each demanding greater organizational readiness, technical depth, and governance maturity.

The 180-Day Activation Roadmap

Transitioning to an agentic enterprise requires disciplined sequencing. The first 180 days should establish controlled momentum with demonstrated operational value and durable foundations.



In line with this roadmap, the objective by Day 180 is demonstrable, governed, and economically justified autonomous execution in selected domains—supported by technical coherence, supervisory clarity, and leadership alignment.



Conclusion:

Preparing the Enterprise for Delegated Judgment

Agentic AI is an organizational inflection point, not merely a technology milestone. As autonomous systems increasingly influence pricing, risk adjudication, supply continuity, and capital allocation, oversight must be formal and active. Boards must delegate judgment without diluting ethics or accountability. Autonomy layered onto fragmented data and weak governance magnifies exposure; embedded within coherent controls, it amplifies institutional capability.

Agentic AI is not a destination. It is a capability. The winners of the next decade will not be those who adopt AI first, but those who build the maturity to let AI act responsibly, safely, and at scale.

The question remains:
**Are you truly
Agentic-Ready?**

