RESEARCH REPORT

Agentic AI Framework

A Structured Framework for Enterprise Adoption Derived from Leading Industry Research

Synthesised from six major industry research publications, April 2026

Sources: Anthropic | Accenture | BCG / MIT Sloan | IBM Institute for Business Value | Bain & Company | OpenAI

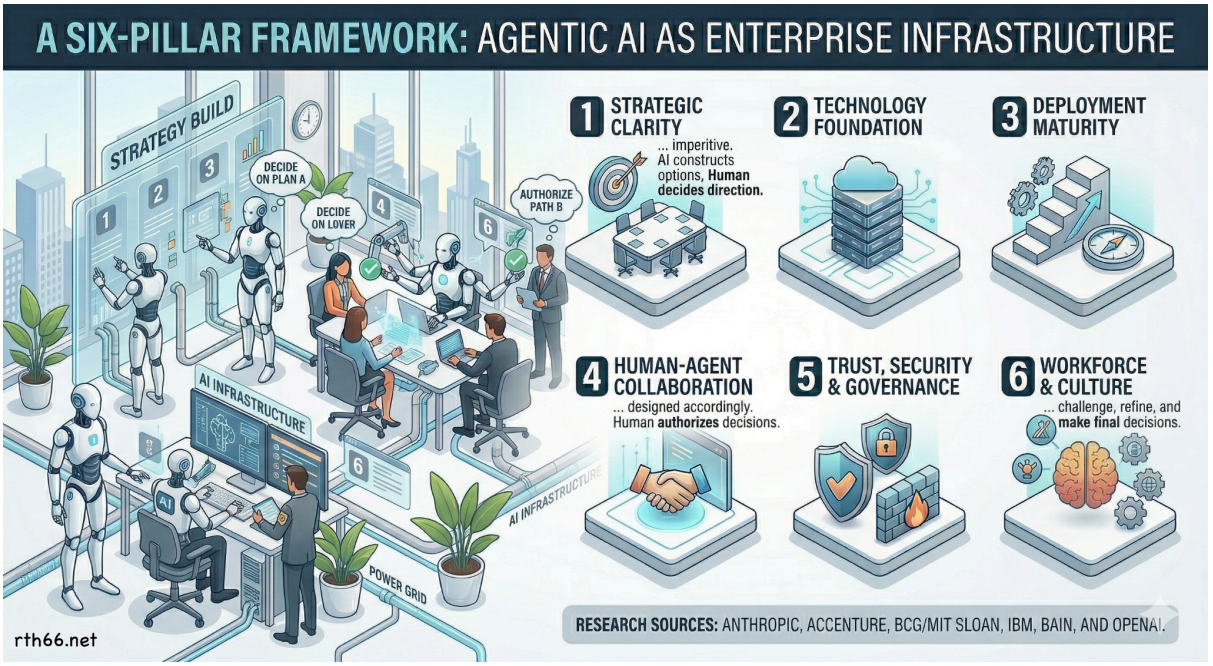

Executive Summary

Agentic AI is no longer a future concept. It is becoming core enterprise infrastructure. Across six major research publications produced by Anthropic, Accenture, BCG/MIT Sloan, IBM, Bain & Company, and OpenAI, a consistent and urgent message emerges: organisations that fail to move beyond pilots and adopt a structured, strategic approach to agentic AI risk falling permanently behind.

This report synthesises the common findings from these six sources and presents them as a practical, six-pillar Agentic AI Framework. The framework is designed to help senior IT leaders, CIOs, and business executives understand what agentic AI is, why it matters now, and how to approach adoption in a structured, responsible, and competitive way.

KEY STATISTIC

80% of organisations report their AI agent investments are already delivering measurable ROI today.

Source: Anthropic, 2026 State of AI Agents Report

KEY STATISTIC

Companies aligning AI, platform, and business strategy achieve 13% more revenue growth and 37% higher operating profit.

Source: Accenture, The New Rules of Platform Strategy in the Age of Agentic AI

KEY STATISTIC

78% of C-Suite executives agree that achieving maximum benefit from agentic AI requires a new operating model.

Source: IBM Institute for Business Value, Agentic AI's Strategic Ascent

The six pillars of the Agentic AI Framework are:

- **Pillar 1: Strategic Clarity** — Define the role of agentic AI at the enterprise level, not the department level.
- **Pillar 2: Technology Foundation** — Build the infrastructure, data, and platform architecture required to support autonomous agents at scale.
- **Pillar 3: Deployment Maturity** — Progress systematically from single-task automation to multi-agent orchestration.
- **Pillar 4: Human-Agent Collaboration** — Redesign workflows, roles, and governance to integrate humans and agents effectively.
- **Pillar 5: Trust, Security, and Governance** — Establish the transparency, oversight, and ethical guardrails that autonomous AI systems require.
- **Pillar 6: Workforce and Culture Transformation** — Prepare people, skills, and organisational culture for an AI-augmented enterprise.

1. What Is Agentic AI?

Agentic AI refers to AI systems that are designed to act autonomously or semi-autonomously to complete goals. Unlike earlier AI models that waited for a human input before producing an output, agentic systems can plan, reason, make decisions, and take multi-step actions on their own. They adapt dynamically to new information and can collaborate with other agents or systems to achieve complex outcomes.

Accenture *defines agentic AI as: "AI systems designed to act autonomously or semi-autonomously to complete goals, often using reasoning, planning and decision-making. These systems can adapt over time and may work collaboratively with humans or other agents."*

BCG/MIT Sloan describe agentic AI as a system that can **plan, act, and learn on its own** — behaving increasingly like an autonomous teammate rather than a tool waiting for instructions. Their research found that 76% of executives already view agentic AI as more like a coworker than a tool.

This dual nature — part tool, part coworker — is what makes agentic AI fundamentally different from all prior enterprise technology. It creates new challenges for governance, investment models, workforce planning, and process design, none of which existing management frameworks were designed to handle.

1.1 The Four Levels of Agentic AI Capability

Bain & Company identify four levels of agentic AI capability that organisations progress through:

Level	Description	Example
Level 1	LLM-powered information retrieval agents	Knowledge assistants, copilots
Level 2	Single-task agentic workflows	Task-doers with self-contained action loops
Level 3	Cross-system agentic workflow orchestration	Complex workflow execution with supervised agents
Level 4	Multi-agent constellations	Any-to-any agent, loosely coupled collaborative agents

Source: Bain & Company, Technology Report 2025 — State of the Art of Agentic AI Transformation

Bain's research indicates that most tech-forward enterprises have successfully scaled Level 1 in 2023-2024. The current competitive frontier is Levels 2 and 3, while Level 4 multi-agent constellations remain emerging.

2. Why the Urgency? The Competitive Reality

Every source examined for this report issues the same warning: the window for first-mover advantage is narrowing, and organisations still in pilot mode are falling dangerously behind.

2.1 Adoption Is Accelerating Rapidly

Agentic AI is spreading faster than any previous enterprise technology wave. BCG/MIT Sloan research shows that traditional AI reached 72% adoption over eight years, while generative AI reached 70% adoption in three years. Agentic AI has already reached 35% adoption in just two years, with another 44% of organisations planning to deploy it soon.

Vendors are accelerating this further by embedding agentic capabilities directly into existing enterprise platforms — SAP, Salesforce, ServiceNow, Workday, Microsoft, Oracle and others all launched agentic features in 2025. This means many organisations are already implementing agentic AI before they have a strategic framework in place.

2.2 Leaders Are Already Pulling Ahead

The performance gap between AI leaders and laggards is widening rapidly and may become permanent.

BAIN RESEARCH FINDING

AI leaders have moved from pilots to profits, delivering 10% to 25% EBITDA gains by scaling AI across core workflows.

Source: Bain & Company, Technology Report 2025

IBM RESEARCH FINDING

Organisations that excel in three key AI adoption areas (cybersecurity integration, ethics embedding, and workflow-specific models) are 32 times more likely to achieve top-tier business performance.

Source: IBM Institute for Business Value, Agentic AI's Strategic Ascent

OPENAI RESEARCH FINDING

Enterprise ChatGPT message volume grew 8x year-over-year; reasoning token consumption per organisation increased 320x — demonstrating deepening, not just wider, AI integration.

Source: OpenAI, The State of Enterprise AI 2025 Report

2.3 The Gap Between Aspiration and Execution Is the Real Risk

IBM's research surfaces a critical distinction that explains why many organisations are not yet seeing transformational returns. There are two distinct groups:

- Process-focused organisations: Laser-focused on optimising existing workflows. Technically proficient but not yet transformational. The majority of AI investment (78%) falls into this category.
- Transformation-driven organisations: Pursuing a dual mandate — improving existing workflows AND creating net-new capabilities. This group achieves measurably greater results across every business KPI.

The core message from IBM is that using agentic AI only to do existing work faster is the wrong ambition. The question transformation-driven leaders ask is not 'How do we do this 20% faster?' but 'What becomes possible when systems can make decisions autonomously?'

3. The Agentic AI Framework: Six Pillars

The following framework distils the common themes, priorities, and recommendations from all six source documents into a structured approach for enterprise leaders. Each pillar is supported by consistent evidence across multiple sources.

#	Pillar	Description	Sources
1	**Strategic Clarity**	Define the enterprise-wide role of agentic AI. Align AI, platform, and business strategy. Commit to transformation, not just optimisation.	*Accenture, BCG/MIT, IBM*
2	**Technology Foundation**	Modernise the data environment and platform architecture. Build interoperable, cloud-native, AI-ready infrastructure.	*Accenture, Bain, IBM*
3	**Deployment Maturity**	Progress from single-task agents to cross-system orchestration and multi-agent collaboration. Follow a proven maturity model.	*Bain, Anthropic, OpenAI*
4	**Human-Agent Collaboration**	Redesign workflows, roles, and decision rights around the combination of human judgment and agent autonomy.	*BCG/MIT, IBM, Accenture*
5	**Trust, Security & Governance**	Embed transparency, explainability, ethics, and cybersecurity into every agentic AI deployment from the ground up.	*IBM, BCG/MIT, Accenture*
6	**Workforce & Culture**	Reskill employees, create new AI-native roles, manage change proactively, and build an AI-fluent culture.	*IBM, Accenture, BCG/MIT*

Pillar 1: Strategic Clarity

The most consistent finding across all six research reports is that agentic AI must be treated as a strategic imperative, not a technology initiative. Organisations that deploy AI agents without a clear enterprise strategy produce fragmented results, duplicated efforts, and low returns.

Accenture's research found that only 18% of companies describe their AI, platform, and business strategies as fully aligned. Yet those that do achieve alignment **double their revenue growth and increase operating profit by up to 37%** compared to peers.

BCG/MIT Sloan reinforce this, finding that competitive benefits from agentic AI will come not from early access to technology but from superior organisational design around it. The technology itself is becoming commoditised. Strategy is the differentiator.

What strategic clarity requires:

- Set ambitious transformation goals based on top-down diagnostics, not bottom-up pilots.

- Charge general managers (not only the CIO or CTO) with accountability for meeting AI-driven business targets.
- Define the enterprise North Star: moving from process optimisation to net-new capability creation.
- Treat agentic AI as enterprise-wide infrastructure, not a department-specific tool.
- Develop new KPIs to monitor agentic AI's impact on business targets — not just traditional efficiency metrics.

Supporting sources: Accenture (2025), BCG/MIT Sloan (2025), IBM IBV (2025)

Pillar 2: Technology Foundation

Every source report identifies data quality, platform architecture, and infrastructure readiness as essential prerequisites for agentic AI. Without a solid technology foundation, agents cannot function reliably or at scale.

IBM's research identifies the top three data-related barriers to agentic AI implementation as: data privacy and security (65%), data integration complexity (60%), and data quality issues (56%). These are not IT problems — they are strategic barriers requiring enterprise-wide solutions.

Accenture's research shows that platform modernisation is the foundation for AI scale. Organisations that consolidate legacy systems, standardise data, and move to modular, cloud-based architectures create the conditions for scale. Fragmented operations give way to connected systems. Manual work becomes automated flow.

Technology foundation priorities:

- Modernise the digital core: consolidate legacy systems, standardise data models, and move to modular, cloud-native architectures.
- Invest in data quality and governance before agent deployment. Process redesign and data cleanup cannot be a shortcut.
- Build platform-aware infrastructure that allows AI agents to scale seamlessly across enterprise systems.
- Evaluate a hybrid approach: 47% of organisations combine off-the-shelf agents with custom-built components, reflecting that no single approach delivers everything needed.
- Consider workflow-specific small language models (SLMs) for precision, transparency, and speed — rather than relying solely on generic large models.

Supporting sources: Accenture (2025), IBM IBV (2025), Bain & Company (2025), Anthropic (2026)

Pillar 3: Deployment Maturity

No organisation should attempt to jump directly to multi-agent orchestration. The research consistently shows that a phased approach — starting with manageable use cases that demonstrate ROI, then expanding — is the path that produces sustainable results.

Anthropic's 2026 research reveals that more than half of organisations (57%) now deploy agents for multi-stage workflows, and 81% plan to tackle more complex use cases in 2026. The top-impact agentic use cases today include: software development (57%), customer service (55%), data analysis and report generation (60%), and internal process automation (48%).

OpenAI's 2025 Enterprise Report documents that enterprise workers report saving 40-60 minutes per day through AI use, with 75% reporting they can now complete tasks they previously could not perform — indicating both efficiency gains and capability expansion.

Deployment maturity principles:

- Start with high-volume, repetitive workflows with clear performance metrics — these provide measurable ROI fastest.
- Prioritise use cases where agents can either amplify expert judgment (data analysis, reporting) or eliminate low-value work (process automation), rather than simply digitising existing manual processes.
- Follow the four-level maturity model: retrieval agents first, then single-task agents, then cross-system orchestration, then multi-agent constellations.
- Maintain an architectural North Star but build pragmatically: fit-for-purpose, domain-specific, human-in-the-loop solutions for the foreseeable future.
- Measure ROI rigorously and compound: organisations that are seeing the most value deploy agents across multiple use cases and build institutional knowledge over time.

Supporting sources: Anthropic (2026), Bain & Company (2025), OpenAI (2025)

Pillar 4: Human-Agent Collaboration

One of the most critical and consistently underestimated aspects of agentic AI adoption is the redesign of how humans and agents work together. The research is clear: agentic AI is not a replacement strategy. It is a collaboration strategy.

BCG/MIT Sloan identify four key strategic tensions that every organisation must navigate when adopting agentic AI:

- Scalability versus adaptability: Tools scale predictably; workers adapt dynamically. Agentic AI does both simultaneously — requiring new organisational design principles.
- Experience versus expediency: When is the right time to invest in agentic systems and how should those investments be made?
- Supervision versus autonomy: Traditional oversight models assume either full human control or complete automation. Agentic AI requires a spectrum.
- Retrofit versus reengineer: When, and by how much, should organisations change existing processes?

Accenture frames this as the need to define the interplay between three actors: humans, platforms, and agents. The most important questions are: What should be automated? What must remain human? How should people and AI collaborate?

Human-agent collaboration design principles:

- Redesign entire workflows — not just isolated use cases — to integrate human judgment and agent execution optimally.
- Define clear decision rights: which decisions are autonomous, which require human review, and which remain fully human.
- Avoid over-standardisation: it eliminates the adaptive capability that makes agentic AI valuable in edge cases and unexpected situations.
- Create new roles: AI orchestrators, autonomous system auditors, and AI workflow designers are emerging as critical positions.
- Treat agentic AI as a colleague, not just a tool: 76% of BCG/MIT survey respondents already do.

Supporting sources: BCG/MIT Sloan (2025), Accenture (2025), IBM IBV (2025)

Pillar 5: Trust, Security, and Governance

As agentic AI systems take on more autonomous decision-making, trust becomes the most important factor in achieving scale. IBM's research identifies lack of visibility into agent decision-making as a significant implementation barrier for 45% of executives. This is not merely a technical limitation — it is a design choice.

IBM identifies three critical drivers that make organisations 32 times more likely to achieve top-tier business performance:

- Integrating cybersecurity into AI initiatives. When agents make decisions independently, a compromised system becomes a weapon. Robust security is not optional.
- Embedding ethics analysis in AI deployments. Without ethical guardrails, agents can perpetuate biases, violate privacy, or make decisions that are logically correct but fundamentally wrong.
- Implementing workflow-specific small language models. Generic AI models struggle with industry-specific jargon and process complexities.

Accenture reinforces that platforms remain essential for reliability, governance, and transaction integrity — they are not replaced by agents, but must be modernised to host embedded intelligence and work seamlessly with agents.

Trust and governance requirements:

- Engineer transparency from the ground up: every automated decision must be auditable and explainable.
- Implement MLOps (Machine Learning Operations) to keep AI models healthy, current, and monitored.
- Use A/B testing and observability logging to track agent decision paths and validate compliance before autonomous execution.
- Set governance standards for agent marketplaces and third-party agents: demand transparency and alignment with strategic KPIs.
- Avoid black-box deployments: the trust deficit is the primary barrier to scaling agentic AI across regulated industries.

Supporting sources: IBM IBV (2025), BCG/MIT Sloan (2025), Accenture (2025)

Pillar 6: Workforce and Culture Transformation

Technology is only half the challenge. Every research source identifies workforce readiness, cultural change, and human capability development as make-or-break factors for agentic AI transformation.

IBM's research finds that 47% of organisations cite inadequate employee skills as a barrier to implementation. Yet the most successful companies are making a strategic bet on their people: 79% of leaders believe they need to protect and value human critical thinking — the skill that will differentiate organisations as algorithms become commoditised.

Accenture's research identifies culture as a distinct transformation priority: organisations must equip people with trust, AI fluency, and the right mindset to lead in an AI-driven enterprise. Importantly, the goal is not replacement but reinvention — integrating people, platforms, and intelligence into a single adaptive architecture.

Workforce transformation priorities:

- Pair AI rollouts with intentional change management: transparent communication, early employee involvement in workflow design, and clear reskilling pathways.
- Build towards AI, not against it: position AI as something built with the workforce, not done to it.
- Create net-new roles: AI transformation executives, AI workflow and experience orchestrators, AI cybersecurity monitors, AI model and data governance leads, and ethical guidelines auditors.
- Develop human critical thinking as a core organisational capability — the ability to challenge, refine, and override autonomous AI decisions and articulate why.
- Measure change readiness and employee sentiment as formal KPIs alongside technical AI performance metrics.

Supporting sources: IBM IBV (2025), Accenture (2025), BCG/MIT Sloan (2025)

4. Cross-Cutting Themes

Beyond the six pillars, the six research reports surface several themes that cut across all dimensions of agentic AI adoption:

4.1 From Optimisation to Transformation

The most important strategic distinction in the literature is between using agentic AI to do existing work faster versus using it to do entirely new things. IBM defines these as 'process-focused' versus 'transformation-driven' organisations. The data is unambiguous: transformation-driven organisations achieve significantly better results across every measurable KPI, including efficiency, revenue growth, customer experience, employee sentiment, and operating margin.

4.2 Data Quality Is Not Optional

Bain, IBM, and Accenture all identify data readiness as the most foundational prerequisite for agentic AI. Process redesign and data cleanup are the most important aspects of any AI transformation — there is no shortcut. Every day an organisation delays this work is another day it falls further behind.

4.3 Architecture Requires Pragmatism, Not Purity

Multiple sources caution against waiting for a perfect architectural solution before proceeding. The pace of AI innovation means a purist approach will not meet the moment. Organisations should maintain an architectural North Star while building pragmatically — using fit-for-purpose, domain-specific, and human-in-the-loop solutions in parallel.

4.4 The Human-Agent Balance Is Dynamic

BCG/MIT Sloan find that organisations must navigate a spectrum of human oversight — from full human control to degrees of automation — rather than choosing between one or the other. Agentic AI systems require supervision like workers do but scale like tools do. New management frameworks, decision rights, and oversight protocols are required.

4.5 Measurement Frameworks Must Evolve

Traditional ROI models, depreciation schedules, and performance metrics are inadequate for agentic AI. IBM finds that only 42% of process-focused organisations have developed new KPIs to monitor agent impact, compared to nearly half of transformation-driven organisations. Leaders need new measurement frameworks that capture continuous learning, emergent capabilities, and the full value of an integrated AI ecosystem.

5. Action Guide for Senior IT Leaders

Based on the synthesised findings from all six source documents, the following actions are recommended for CIOs, senior IT leaders, and IT advisory professionals in the DACH and European market:

Action	What to do	Sources
Assess your current agentic AI level	Use the four-level maturity model (Bain) to honestly assess where your organisation sits today. Identify the biggest gaps and prioritise.	*Bain*
Align strategy before deploying tools	Ensure AI, platform, and business strategy are aligned at board level. Fragmented strategy produces fragmented results.	*Accenture, IBM*
Start with data cleanup	Identify and address data quality, integration, and governance issues before deploying agents. This is the most important prerequisite.	*Bain, IBM*
Choose the right use cases first	Prioritise high-volume, data-rich workflows with clear metrics: software development, customer service, data analysis, and process automation.	*Anthropic, OpenAI*
Build trust mechanisms into every deployment	Implement observability, MLOps, and auditability from day one. Address the trust deficit proactively.	*IBM, BCG/MIT*
Redesign workflows, not just automate them	The biggest gains come from redesigning entire workflows around human-agent collaboration, not layering agents onto broken processes.	*Bain, BCG/MIT*
Create new roles now	Begin identifying and developing AI orchestrators, governance leads, and ethics auditors. Do not wait for the technology to mature first.	*IBM, Accenture*
Develop new KPIs	Define what success looks like for agentic AI beyond traditional efficiency metrics. Measure net-new capabilities and transformation outcomes.	*IBM*
Integrate cybersecurity from the start	Autonomous agents that can make decisions create new attack surfaces. Cybersecurity is not a later-stage consideration.	*IBM*
Communicate change proactively	Transparency about how roles will evolve, early employee involvement, and clear reskilling pathways reduce cultural resistance.	*IBM, Accenture*

6. Source Documents

This research report is based exclusively on the following six publications. All findings, statistics, and recommendations are attributed to these original sources throughout the document.

Abbreviation	Full Reference	Research Scope
Anthropic (2026)	The 2026 State of AI Agents Report: How Enterprises Are Building and Deploying AI in Production. Anthropic / Material Research. Late 2025, surveying 500+ technical leaders in the United States.	Enterprise AI agent adoption patterns, use cases, ROI, and deployment trends across company sizes and industries.
Accenture (2025)	The New Rules of Platform Strategy in the Age of Agentic AI: Five Priorities to Help Companies Align People, Platforms and Intelligence. Accenture Strategy. Survey of 1,000+ executives across 12 countries and 10 industries.	Platform strategy evolution, AI-business-platform alignment, five priorities for enterprise platform reinvention.
BCG/MIT (2025)	The Emerging Agentic Enterprise: How Leaders Must Navigate a New Age of AI. MIT Sloan Management Review & Boston Consulting Group. Global survey of 2,102 respondents across 21 industries and 116 countries, plus 11 executive interviews.	Strategic tensions in agentic AI adoption, governance, workforce planning, human-agent collaboration frameworks.
IBM IBV (2025)	Agentic AI's Strategic Ascent: Shifting Operations from Incremental Gains to Net-New Impact. IBM Institute for Business Value. Survey of 800 C-suite executives across 20 countries and 19 industries.	Operating model redesign, trust and transparency in autonomous AI, workforce evolution, and performance differentiation.
Bain (2025)	Technology Report 2025: AI Leaders Are Extending Their Edge. Bain & Company. Sixth annual Technology Report based on client work and market research across global technology markets.	Agentic AI maturity levels, SaaS disruption, AI infrastructure, ROI patterns, and building the foundation for agentic AI.
OpenAI (2025)	The State of Enterprise AI: 2025 Report. OpenAI. Based on de-identified, aggregated usage data from 1 million+ business customers and a survey of 9,000 workers across 100 enterprises.	Enterprise AI usage patterns, productivity impact, workflow integration depth, and the growing gap between AI leaders and laggards.

Prepared: April 2026