LATENTTI

DOCUMENT

Future of Enterprise Architecture 2026–2031 — Technical Synthesis

AUTHOR LATENTTI

DATE 2026-05-02

Draw the line before it forms by accident.

Contents

Future of Enterprise Architecture 2026–2031

Technical synthesis — Part A (sections 1–5b)

For senior decision-makers responsible for enterprise architecture — CTOs, CDOs, Chief Architects, transformation leaders, and accountable boards. Specialist work routed through cross-examination, bias counters, and Chain-of-Verification. Confidence: `[certain]`, `[likely]`, `[speculative]`, `[unknown]`. Source tiers: [T1] primary, [T2] reputable secondary, [T3] vendor commentary.

---

1. Central thesis

The 2026–2031 transition is not one paradigm replacing another. It is **three loosely coupled threads arriving simultaneously, with directional asymmetric pull between them**. Agentic orchestration moves the integration substrate from setup.exe-paradigm contracts to prompt-paradigm capability calls (MCP, A2A, AG-UI, A2UI). Knowledge stratification replaces single-corpus enterprise search with multi-tier personal → team → function → corporate → cross-org → public architectures with provenance, decay, governance gates. Determinism routing turns deterministic vs. probabilistic execution into an explicit, governed engineering decision. Bundling these as "EA 3.0" is rhetorically convenient and analytically lossy [AI Architect dissent, accepted]. The protocol thread pulls the other two along at roughly half a curve's lag; identity work runs on its own clock and lags by 3–5 years [likely; cross-exam Pair 4].

The constraint deciding the trajectory is **not capability — it is the gap between adoption and integration**. Stack Overflow 2025: 84%+ developer AI use; JetBrains 2025: only 44% workflow integration; trust fell 70%+ → 60% [T1, T2]. McKinsey 2025: 88% organisational AI use, ~33% scaled [T2; CoVe Claim 2]. Forrester predicts 75% of firms will fail at building agentic architectures **independently, without vendor partnerships and mature governance** [T2; CoVe Claim 3]. The defensible middle is **bifurcation as the modal trajectory** — frontier firms compounding integration depth while the median Fortune 500 stays in surface adoption, with Nordic public-sector cooperative-platform programmes running as a structurally distinct third overlay [T1]. Probabilities are bands: Scenario C (Bifurcation) **40–50%**, Scenario A (Compounding) **25–35%**, Scenario B (Plateau) **25–35%** [Trend Analyst §5; bias-counter §1; CoVe Claim 7]. Bifurcation is modal; not more probable than the alternatives combined. Inside

C the bifurcation is **per-thread**: Thread A consolidates near-uniformly; Threads B and C bifurcate sharply on organisational discipline; identity bifurcation lags B and C by 2–3 years and looks more like uniform identity-lag in 2031 than visible bifurcation [Cross-exam Pair 5].

The load-bearing through-line is the **cognitive turn**. Across all five populations the shape repeats: tool/output self-concepts ("I write COBOL"; "I run servers"; "I draw the integration map"; "I make the call") face cliff-edged compression. Discipline self-concepts ("I keep the platform safe"; "I steward the data layer"; "I architect the user-system contract"; "I am the routing-policy designer") cross the shift with craft intact and often elevated. **The cliff is on tool/output identity. The continuity is on discipline identity** [HR Advisor; baseline §D7].

Three operating-model postures: Frontier (build the platform, accept the failure rate), Fast-follower (wait for protocol standardisation and hyperscaler packaging, spend on governance and change), Specialist (capability provider). **Fast-follower is modal**; most firms hedge unevenly [Platform Economist]. The 2031 picture: deterministic where it matters, agentic where it pays, continuous compliance binding the two — and a workforce whose identity language has either shifted or hollowed out.

---

2. The change reframed: three threads, not one paradigm

The "EA 3.0" frame describes a real shift but bundles threads that mature on different curves with different uncertainty profiles. Naming them separately is more honest and makes the case for change stronger. The capability-trajectory dependency map (§8 in Part B) distinguishes claims that hold under capability plateau from claims dependent on continued compounding.

Thread A — Agentic orchestration (the protocol substrate)

The clearest, fastest-moving thread. MCP went from Anthropic-only release (Nov 2024) to OpenAI adoption (March 2025), Google DeepMind (post-March 2025), Microsoft Copilot Studio (July 2025; **45M SDK downloads**), AWS support (November 2025; **68M SDK downloads**), and Linux Foundation governance under the Agentic AI Foundation (December 2025) — **a 13-month vendor-neutralisation arc, faster than Kubernetes' roughly three-year arc to CNCF** [T1, T2; CoVe Claims 6, 15]. AG-UI (CopilotKit-stewarded) and A2UI (Google declarative spec) sit earlier on the same arc [T1]. This thread carries `[likely → certain]` confidence and holds under both compounding and plateau scenarios.

Thread B — Knowledge stratification (data → memory)

Highest strategic stakes, most timing uncertainty. Glean's $200M ARR / $7.2B valuation after a $150M Series F (June 2025) [T1] proves the corporate-tier knowledge layer is independently valuable; Microsoft Copilot for M365, Google Agentspace, AWS Q, Notion AI build team and function tiers [T1]; agentic memory primitives (Letta, Mem0, Zep) productise the foundation [T1]; LazyGraphRAG drops indexing cost to ~0.1% of full GraphRAG [T1]. Missing in 2026: governance for the boundaries between tiers — escalation, decay, redaction, forgetting, classification, cross-org sharing under purpose limitation [Security Auditor §3]. This is where Forrester's "knowledge curator" lives [T2]. Adoption-curve confidence `[likely]` through 2028; mainstream multi-tier with governed forgetting `[speculative]` for 2031, capability-trajectory-dependent.

Thread C — Determinism routing (engineered choice, not default)

The thread most easily over-claimed and most easily missed. The capability is real today — structured outputs, function calling, AEGIS-style guardrails, OPA Rego / AWS Cedar policy-as-code, hybrid retrieval-then-rule patterns common in Finnish and Swedish public-sector administrative-law-bound architectures [Security Auditor §1, §7; AI Architect §5]. The thread defines where the other two stop. Latency-, cost-, audit-, and regulator-critical surfaces stay deterministic; high-variance, language-heavy, generative-interface surfaces move to agentic;

the bridge layer is where 2026–2031 engineering interest concentrates. Under plateau, Thread C strengthens — when agents are unreliable the deterministic envelope expands. Confidence `[likely]` on pattern; routing topology `[speculative]` and sector-specific.

Thread A matures fastest (consolidation effectively complete by 2027–2028). Thread B matures unevenly — frontier 2029, median lags 2–4 years. Thread C is permanent engineering discipline. Retire "EA 3.0" in favour of **"three threads arriving simultaneously, on different curves, with directional asymmetric pull, reinforcing each other."**

---

3. The five populations × five-year trajectory

3.0 Cross-population overview

The master matrix below summarises trajectory shape per population. Detail follows in §3.1–§3.5.

POPULATION	NOW (2026)	FORCES	+24MO (2028)	+60MO (2031)	SELF-CONCEPT SHIFT	TECH PARTICIPATION	KNOWLEDGE-TIER ROLE	DETERMINISM-ROUTING ROLE
P1 The firm	200–400 SaaS apps; ~33% scaled AI; <20% have capability inventory [T1, T2]	Composable / MCP-exposed services; capability registry; hyperscaler vs. open-protocol; Forrester 75% fail at independent build [T2]	Capability registry pilots; first agentic process wins; hyperscaler registry GA	Frontier 15–25%: portfolio-of-capabilities; modal: capability-augmented stack [VISION; speculative on cohort size; likely on direction]	"[Vertical] company" → "outcome-delivering organism orchestrating human + agent capability" [HR Advisor §5]	Assistant universal; agentic mainstream 2029–2030; autonomous bounded	Sets policy on tier topology, classification, cross-org sharing, governed forgetting	Sets firm-level policy: where determinism is mandated, where agency is permitted

POPULATION	NOW (2026)	FORCES	+24MO (2028)	+60MO (2031)	SELF-CONCEPT SHIFT	TECH PARTICIPATION	KNOWLEDGE-TIER ROLE	DETERMINISM-ROUTING ROLE
P2 Operational worker	75% gen AI use; 44% integrated; identity = executional competence [T1]	Routine eaten; metrics shift to override-rate + outcome quality; bimodal at cohort level [Operations Manager §2]	Agent composes routine processes; "agent boss" mainstream; outcome metric primary	Median customer-ops team ≤50% of 2025 FTE [BET, likely on shape; speculative on threshold; 45–60% plausible band]	"Process executor" → "process composer / judgment layer / agent boss"; bimodal at cohort, smoother-uneven at individual [Cross-exam Pair 6]	Assistant certain; agentic mainstream 2027–2028; autonomous bounded	Consumes function/ corporate; contributes override patterns to team tier	Lives the routing decision daily

POPULATION	NOW (2026)	FORCES	+24MO (2028)	+60MO (2031)	SELF-CONCEPT SHIFT	TECH PARTICIPATION	KNOWLEDGE-TIER ROLE	DETERMINISM-ROUTING ROLE
P3 Leader-as-leader	Decision-maker, strategist, results-driver; AI as background tool [HR Advisor §2]	Operational decisions enter agent-mediated workflows; "agent boss" upward; sense-maker reframing	Agent-mediated portfolios; hybrid teams; decision rituals reshape	Trust-architect, routing-policy designer; rewritten career architecture in lead-edge firms [likely for leading-edge; speculative for median]	"Decision-maker" → "trust architect, routing-policy designer, sense-maker" — 2028–2029 inflection at leading edge	Assistant certain; agentic delegation rises 2027–2029; autonomous narrow post-2029	Net contributor: rationales become corporate-tier training data	Primary stakeholder of routing policy

POPULATION	NOW (2026)	FORCES	+24MO (2028)	+60MO (2031)	SELF-CONCEPT SHIFT	TECH PARTICIPATION	KNOWLEDGE-TIER ROLE	DETERMINISM-ROUTING ROLE
P4 Technical builder	84–85% AI use; code-author identity; trust 70→60% [T1, T2]	Routine implementation eaten; verification overhead grows; agent integrator + verifier engineer + knowledge architect emerge	Microsoft Agent Framework + LangGraph dominant; specification-as-craft; junior pipeline question acute	Stable headcount, reshaped mix; system-shepherd self-concept dominant for senior tier; junior pipeline re-invented in lead-edge firms [BET]	"Code author" → "intent specifier and validator; system shepherd" — easier for senior, harder for junior	Assistant universal; agentic standard for low-risk surfaces 2028; autonomous narrow but growing	Builds MCP servers per tier, retrieval policy, agentic memory, eval harnesses	Builds the bridge layer

POPULATION	NOW (2026)	FORCES	+24MO (2028)	+60MO (2031)	SELF-CONCEPT SHIFT	TECH PARTICIPATION	KNOWLEDGE-TIER ROLE	DETERMINISM-ROUTING ROLE
P5 Whole-system architect	Integrator/ standards-keeper; classical TOGAF anchor [T1]	Forrester 4 emerging roles fragment classical role; agentic governance + knowledge architecture become core; ARB cycles compress [T2]	Knowledge architect role recognised in analyst taxonomies; capability registry curation displacing static maps	Bifurcation: ecosystem cultivator + 4 specialist sub-roles + classical-EA legacy niche [BET]	"ARB police / model-keeper" → "ecosystem cultivator / routing-policy steward / agentic governance champion" — branching	Assistant universal; agentic core to discipline; autonomous bounded	Governs topology, classification, governance gates	Owns the architectural primitive

POPULATION × CAPABILITY HEATMAP

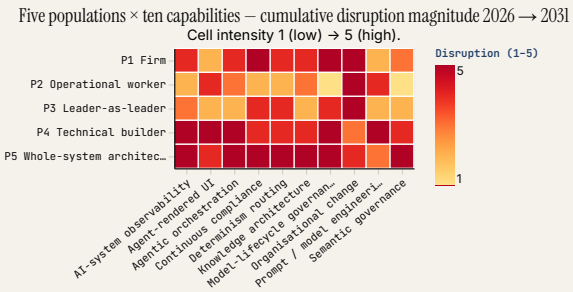

P4 and P5 carry the highest summed-disruption rows. P3 concentrates in organisational change, model-lifecycle governance, determinism routing, continuous compliance — the leadership-craft cluster. P2 is hit on agent-rendered UI plus organisational change.

SHIFT-INTENSITY CURVE

Normalised role-shift intensity 2026 → 2031 per population
Trend Analyst §4 with reference-class adjustment. [likely] on shape; [speculative] on year-on-year values.

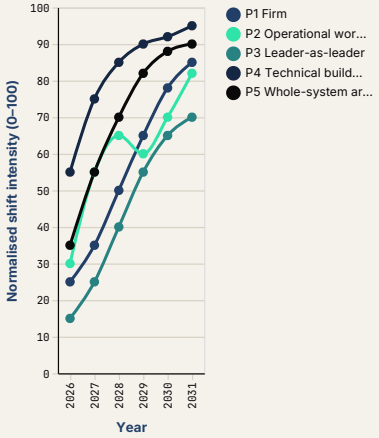

P4 starts highest and steepens earliest; P5 catches P4 by 2031; P1 climbs steadily; P3 trails until the 2028–2029 cognitive turn lands at the leading edge; P2 is bimodal at cohort level — first wave 2026–2028, plateau 2029, second wave 2030–2031.

3.1 Population 1 — The firm

The firm is the unit of analysis, not the application. For two decades the enterprise was legible as a stack of products and capability was conflated with licence count [Platform Economist §1]. Composable design (MACH at design-time, MCP-style addressability at runtime) breaks the conflation; agentic composition turns capability binding into a runtime decision. The firm of 2031 is a portfolio of addressable capabilities and the knowledge that contextualises them.

NOW (2026). Mid-large enterprises run 200–400 SaaS products plus custom layer and 50–150 active integrations [Platform Economist §2.1]. McKinsey 2025: 88% AI use, 33% scaled, 32% expecting headcount decreases concentrated in service operations, customer care, supply chain [T2]. Microsoft WTI 2025: 75% knowledge-worker generative AI use, 33% of leaders considering headcount reductions, 78% considering hiring for new AI roles, "agent boss" coined [T1]. Composable adoption ~10–20% deliberate; the rest run accidental partial-composability. Dominant self-concept: "we are a [vertical] company that uses technology" [HR Advisor §5].

Forces. MCP, A2A, AG-UI, A2UI compress integration economics from "build the integration" to "describe the capability"; integrator role compresses, capability-curator expands. Hyperscaler gravity wants orchestration in one mesh; open-protocol gravity wants vendor-neutral substrate. Most firms pay integration tax to both. Forrester's 75%-fail-at-independent-build [T2] is the empirical floor.

Trajectory NOW → +24mo → +60mo.

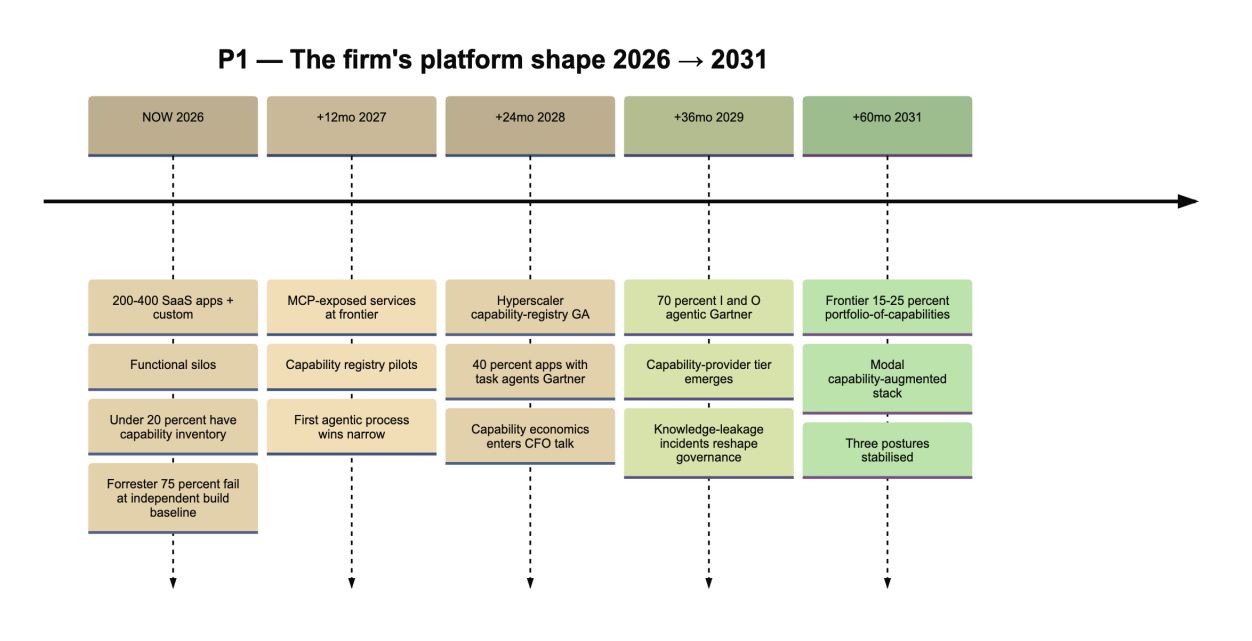

Self-concept shift. "[Vertical] company" → "outcome-delivering organism that orchestrates human + agent capability." The identity locus moves from vertical to orchestration capability. The leading 15–25% complete the shift; the median lags 3–5 years [VISION; speculative on cohort size; likely on direction].

Tech participation. Assistant universal by 2027. Agentic frontier 2027–2028, mainstream 2029–2030. Autonomous bounded but growing — supply-chain rebalancing within guardrails, dynamic pricing within bands, simple-claims auto-adjudication. **Knowledge-tier role.** Sets policy on tier topology, classification, cross-org sharing, governed forgetting. Corporate-tier knowledge is the durable moat when governed; commodity when leaked through ungoverned agent calls or third-party retention defaults [Security Auditor §3]. **Determinism-routing role.** Firm sets policy on which decisions are deterministic (regulated, fiduciary, customer-trust-binding, safety-critical) and which are agentic. Finnish administrative-law (Chancellor of Justice OKV/2213/70/2024, ruling 2025-04-02) and Swedish AI-verkstad anchor the deterministic envelope in Nordic public sector [T1; CoVe Claim 16].

Three postures — Fast-follower modal. Frontier: build the platform, accept Forrester 75% as cost of leadership; suits digital natives and frontier industrial firms. Fast-follower (modal): wait for protocol standardisation and hyperscaler packaging; spend on knowledge stratum, governance, change; suits regulated and asset-heavy firms. Specialist: expose narrow-deep capabilities via MCP and A2A; capture network effects on the capability, not the application. Most hedge unevenly. The 75% fail rate concentrates in firms attempting Frontier without the capital or capability [Platform Economist §5].

Cost recalibration. Vertical regulated B2B SaaS: traditional $20–30M / 3–4 yrs / 20–30 engineers → mature 2027–2029 $5–12M / 18–30 mo / 5–15 engineers; lean $3–7M / 12–24 mo / 3–8 [Platform Economist §4]. Mid-complexity non-regulated B2B SaaS is smaller absolute dollars but the same shape [AI Architect §2.6]; the regulatory delta is the point — do not blend the two [CoVe Claim 9]. Enterprise IT maintenance: integration FTE compresses 50–70% on automatable surfaces; total IT cost compresses 25–35% because regulatory, audit, change-management, vendor-management, and trust functions persist or grow [AI Architect §6.2]. The cost base **changes shape, not size**.

3.2 Population 2 — Operational workers

Operational work is organised today around the SOP: SOP codifies steps; worker executes; supervisor measures adherence; metrics reward throughput-on-steps [Operations Manager §1]. The transitional state is "SOPs executed with AI assistants" — worker still owns steps, AI drafts and summarises, throughput up, identity unchanged. This is the 75%-using / 44%-integrated stuck middle. Destination: outcome-based agentic process where humans set outcomes and constraints, agents compose the process, metrics measure outcome and exception quality.

NOW. Self-concept: "I do my job — execute the process, hit the metric." Self-worth tied to executional competence. Pain points: tool sprawl; task-level help, workflow-level friction; distrust on customer-facing and compliance surfaces [Operations Manager §2].

Trajectory.

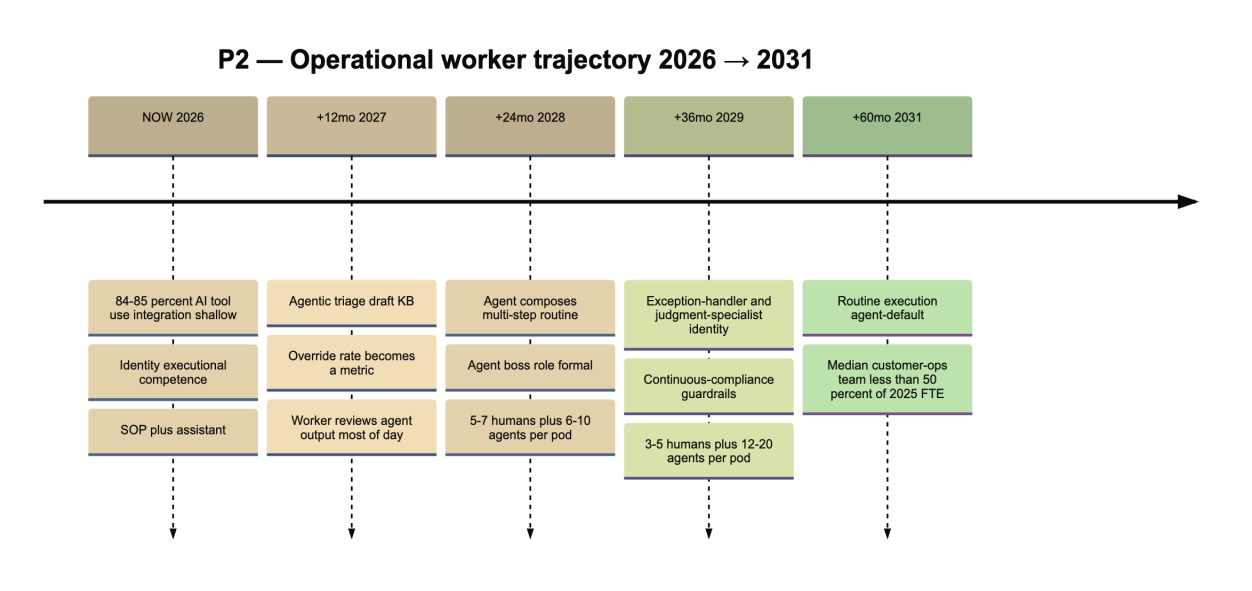

Self-concept shift — bimodal at cohort, smoother-uneven at individual. The worker moves from process executor to one of three new self-concepts: exception handler (cases the agent flags outside confidence band); judgment specialist / customer-trust holder (relationship, regulatory edge cases, complex multi-party deals); or agent boss / orchestrator (building, delegating to, managing, evaluating agents) [HR Advisor §3, Operations Manager §2]. **The hardest cognitive turn is that identity has been tied to the part of the job that compresses most. The cliff is between cohorts, not within an individual's experience** [Cross-exam Pair 6]: ~30–50% navigate to a new self-concept; 20–35% leave the operational track; a residual cohort persists in narrowed SOP niches in regulated or trust-sensitive contexts [Operations Manager §7 VISION].

Cost recalibration — the 45–60% FTE compression band. Mid-size SaaS customer ops (8,000–15,000 customers): traditional 40–80 FTE / ~$4–8M → mature 2030+ 15–30 FTE + platform / ~$2–4M. Mid-size enterprise back-office finance: traditional 20–40 FTE / ~$2–4M → mature 6–15 FTE + platform / ~$1–2M. Confidence: `[BET, likely on shape; speculative on threshold; 45–60% plausible; 75%+ possible under faster autonomous-tier deployment]` [bias-counter §3, §5; CoVe Claim 8]. **What does not compress:** complex judgment, escalation, account-relationship, regulatory escalation, organisational change cost (process redesign, identity-shift coaching, severance, governance redesign — the most-underestimated business-case line item) [Operations Manager §3].

Knowledge-tier role. Consumes function and corporate tiers; **contributes override patterns and exception decisions back to team tier** — the tier most missing in 2026 [Operations Manager §2]. **Determinism-routing role.** Lives the routing decision daily — wrong toward over-automation produces bored gatekeepers; wrong toward under-automation drags agents into assistant-only mode and the integration gap persists.

3.3 Population 3 — Leaders-as-leaders

The population where the cognitive turn is hardest to see from inside, longest in elapsed time, and most consequential — leaders set the routing policy that determines where every other population lands [HR Advisor §2].

NOW. Self-concept: decision-maker, strategist, people-manager, results-driver — four facets that have anchored leadership identity for thirty years [HBR / MIT Sloan 2025, T2]. The adoption-vs-integration gap at leadership level is **likely worse** than at IC level — leaders use AI as meeting-prep, email-drafting, dashboard-interpretation, rarely as a thinking partner. Leaders think about AI as a labour-allocation lever, not yet with AI as an instrument of their own craft.

Forces. Agentic systems eating the operational-decision layer; senior leader becomes boss-of-bosses-of-agents. Board pressure: AI strategy and governance without the leader being a technologist. Identity threat from sense-maker reframing — for leaders anchored in making the call, the reframing reads as demotion. Leader meeting notes, decision rationales, strategy artefacts become training data for next-generation agents [HR Advisor §2].

Trajectory.

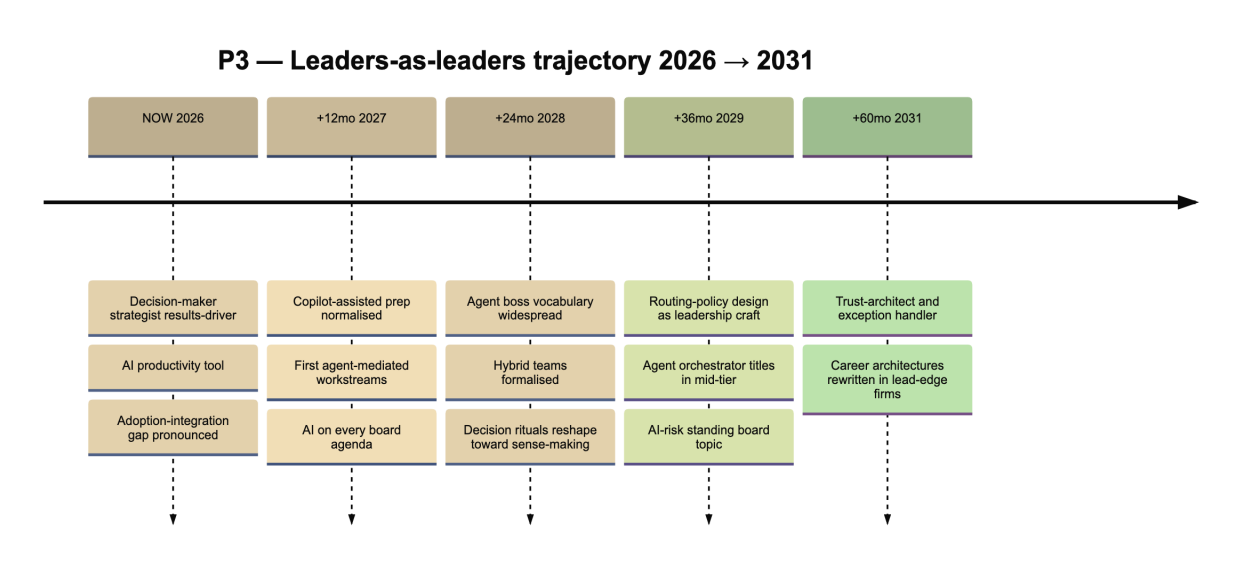

The 2028–2029 cognitive turn — "agent boss" mainstreaming → trust architect identity. The arc has a real inflection where the additive frame ("I lead a team that uses AI") gives way to substantive ("I lead a hybrid team — humans and agents"). The deeper turn to "orchestrator of outcomes" is uncertain at the median by 2031; the leading edge will be there [likely for leading-edge; speculative for median]. Reference class: digital-transformation-leader identity — discourse ~2014, leading-edge mainstream 2018–2020, still incomplete at median in 2024. The agentic turn follows the same shape: leading-edge 2028–2029, median 2030–2032 [bias-counter §6]. The efficiency leader vs. resilience leader tension is not resolvable by tooling. The mature 2031 self-concept holds both: I am the routing-policy designer; I trade efficiency and resilience consciously.

Tech participation. Assistant certain. Agentic delegation rises 2027–2029. Autonomous selectively post-2029 — micro-pricing within bands, capacity scaling within SLA, vendor auto-renewals below threshold. **Knowledge-tier role.** Net contributor to corporate tier. Personal-tier knowledge becomes more strategic (mental models, judgment heuristics, network-of-trust) and less operational. **Determinism-routing role.** Primary stakeholder. Sets policy on where determinism wins (compliance, financial controls), where agency wins (creative, exception handling, customer-trust moments), and where humans remain in the loop on judgment grounds (hiring, firing, ethics calls, escalations above threshold). A leadership competency by 2029 and likely a board-reportable one by 2031.

3.4 Population 4 — Technical system-builders

P4 is the engineering organisation — software developers, platform engineers, data engineers, ML engineers, AI engineers, SREs. The five-year story is **not "developers replaced"**; it is "the role boundary between developers, platform, data, ML compresses; job content shifts from authoring to specifying-and-validating; a new sub-role — agent integrator / knowledge architect / verifier engineer — emerges from inside the population" [AI Architect §2].

NOW. Self-concept: "I'm a software engineer. I write code." Self-worth anchored in craft [Stack Overflow + JetBrains 2025, T1, T2]. Trust decline (70% → 60%) and 44% workflow-integration are the leading edge of identity friction [HR Advisor §4]. Copilot RCT 55.8% faster on isolated tasks (95% CI 21–89%; arXiv 2302.06590; CACM 2024) [T1] but no statistically significant change in commit-based output in longitudinal studies (arXiv 2509.20353) [T1] — both real and not contradictory: AI compresses individual tasks dramatically; AI does not compress multi-month system-build proportionally because the binding constraint shifts from typing to specification and validation [AI Architect §2.1; CoVe Claim 13].

Trajectory.

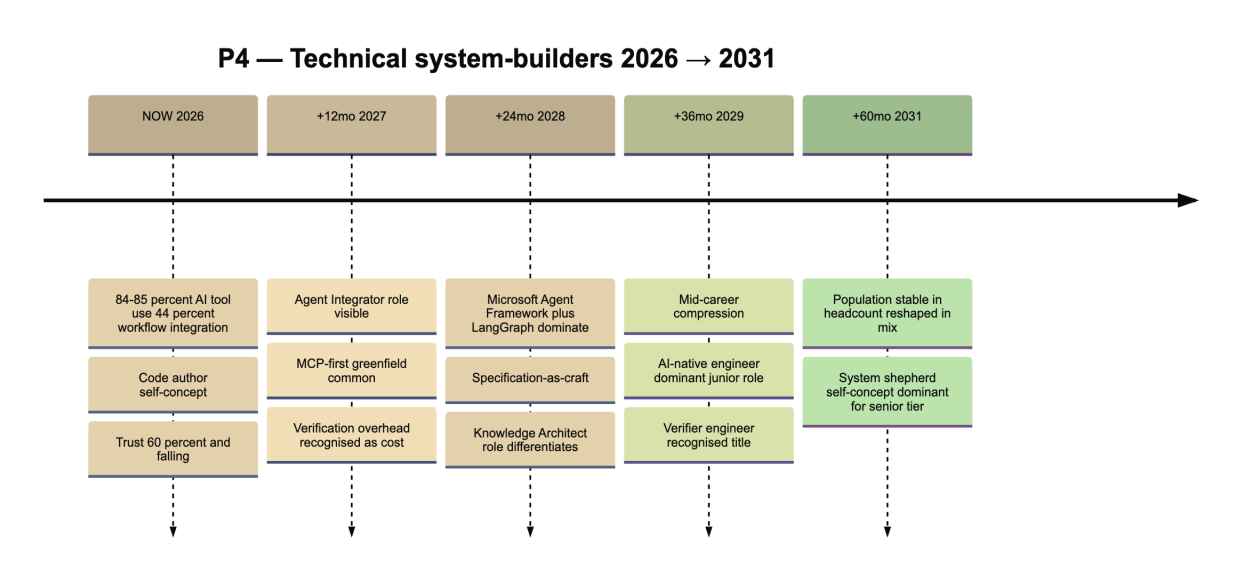

Self-concept shift — easier for senior, harder for junior. "Code author" → "intent specifier and validator; system shepherd; engineer of system judgment." This is a return to the Brooks / Dijkstra / Parnas discipline — the thinking part of software engineering, not the typing part [HR Advisor §4]. The cognitive turn for senior engineers is therefore easier than for juniors.

Junior pipeline crisis. The classical path to senior — write foundation code, debug, internalise through repeated contact, take on more complex problems — breaks when the agent writes the foundation. **How does someone become senior if they never wrote the junior code?** Unresolved. GitHub Octoverse: 80% of new developers use Copilot in their first week — first-

week identity is already AI-paired [T1]. By 2031, lead-edge firms will have re-invented apprenticeship through scaffolded sandboxes restricting agent-layer access during specific learning windows, compounding a 5+ year senior-engineer pipeline advantage [BET; HR Advisor §10 / AI Architect §8].

Cost recalibration — building a SaaS platform. Mid-complexity non-regulated B2B SaaS, 0 → $1M ARR: traditional 6–10 engineers × 12–18 mo ≈ $2.0–3.5M; 2026 baseline 3–5 × 9–12 mo ≈ $0.9–1.8M; 2029–2031 plausible 1–3 × 6–9 mo ≈ $0.3–0.8M [AI Architect §2.6]. Vertical regulated B2B SaaS runs an order of magnitude larger because of regulatory posture and trust with regulated buyers [Platform Economist §4]. **What does not compress:** customer discovery, problem framing, market validation (~30–40% of total time, unchanged); trust and verification work (grows in absolute terms); regulatory work (DORA, GDPR, EU AI Act, SOC 2, HIPAA — adds 15–25% overhead); sales, GTM, customer change, brand, design. **The engineering organisation is not 5–10× smaller in 2031; it is modestly smaller (1.3–2× per unit of feature output) and substantially reshaped in skill mix.**

Tech participation. Assistant universal (Gartner: 90% by 2028 [T2]). Agentic 2026–2029 growth bound by verification and trust, not capability. Autonomous narrow but growing — DevOps remediation, infrastructure scaling, log triage by 2026; broader business-process loops where cost of error is bounded and audit trails airtight by 2031 [AI Architect §2.5]. **Knowledge-tier role.** Builds the architecture: MCP servers per tier, retrieval policy engines, agentic memory infrastructure, eval harnesses, decay machinery [AI Architect §3.3, §3.5]. **Determinism-routing role.** Implements routing decisions; builds the bridge layer (structured outputs, AEGIS-style guardrails, Semantic Kernel policy hooks, hybrid retrieval-then-rule, fallback paths) [AI Architect §5.3].

3.5 Population 5 — Whole-system architects

P5 is the engagement's centre of gravity. The discipline is being extended with sub-domains that did not exist in 2024 — agentic governance, knowledge architecture, determinism routing, model-lifecycle governance, AI-system observability. Forrester names four emerging EA roles (value mappers, digital twin strategists, knowledge curators, agentic governance champions) [T2; CoVe Claim 4]. These are four self-concepts classical TOGAF practitioners must choose between, not four job titles [HR Advisor §6].

NOW. Dominant self-concept: "integrator / standards-keeper / planner." Activist variants: "digital transformation agent" (Kotusev 2024); critical variants: "whole-enterprise architect, IT-frame skeptic" (Graves) [T1]. TOGAF 10 and ArchiMate 3.2 have no agentic-AI module; the framework body moves slower than practitioner reality [T1]. Open Group's 2024 Architecture Roles and Skills Series Guide is the first framework-body acknowledgment of multiple identity tracks.

Forces. Forrester's four emerging roles fragment the classical role; agentic governance and knowledge architecture become core EA capabilities; ARB-gate cycles for routine change replaced by continuous compliance and policy-as-code; standards-as-prose deprecated; capability inventory becomes the operational substrate agents call against [Platform Economist §6, AI Architect §4.2, Security Auditor §6].

Trajectory.

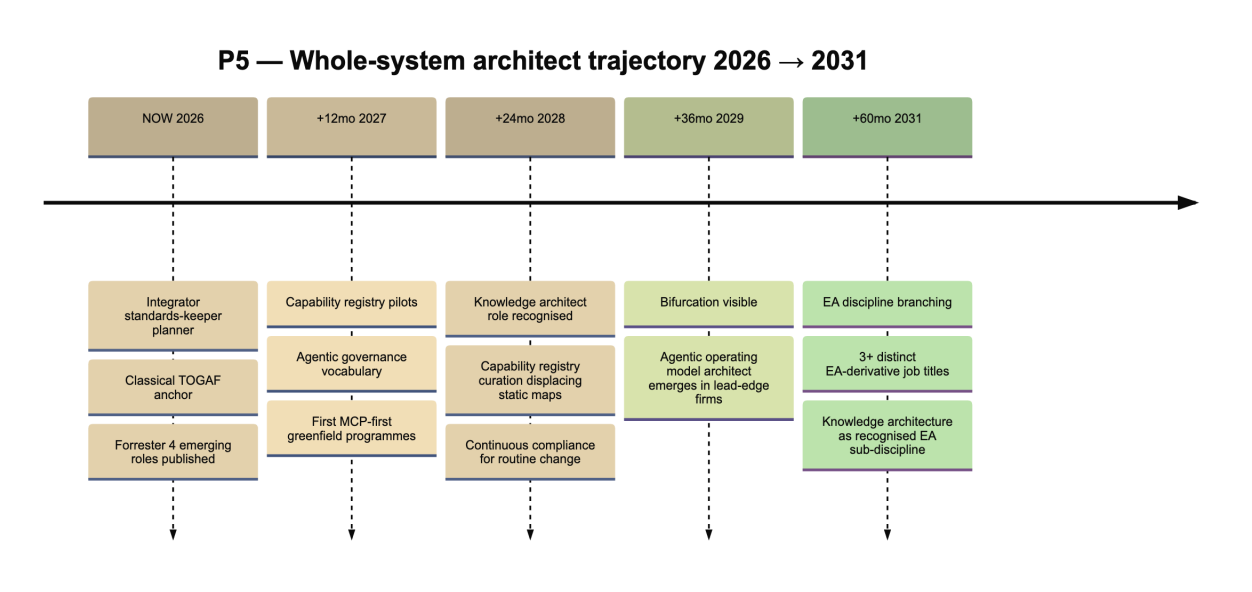

Six-hats applied to the EA role.

- **Facts (white).** Forrester four emerging EA roles published 2025 [T2]. TOGAF 10 in AsciiDoc-on-Git since 2025 [T1]. arXiv 2510.22003 systematic lit review of GenAI's impact on EA in agile environments (Oct 2025, 33 studies) [T1]. Forrester EAMS Landscape Q4 2025 [T2]. MCP Linux Foundation governance Dec 2025 [T1]. Finnish EA salary 9,000–13,000 EUR/month [T2].
- **Risks (black).** EA dissolves into CTO / COO mandates with EA-as-team rather than EA-as-discipline [HR Advisor §6 Path D, speculative]. Mid-career EAs without portable abstractions are most exposed. Forrester 75% is the operational risk on Frontier-mode programmes. Knowledge-leakage incidents 2027–2029 reshape governance under regulator pressure [Platform Economist BET-4]. AI-system security maturity lags deployment by 3–4 years; mid-curve incidents 2027–2029 with probability 40–55% [Security Auditor §11; bias-counter §5; Cross-exam Pair 7].
- **Benefits (yellow).** EA's craft of identifying, naming, bounding capabilities becomes operationally load-bearing rather than a shelved artefact. Senior EAs fluent in capability-economics, knowledge-governance, and agentic-governance command a scarcity premium. Knowledge architecture emerges as a recognised EA sub-discipline. The

discipline is elevated, not eliminated, for practitioners holding the discipline-level self-concept.

- **Creativity (green).** New EA artefacts: agentic operating-model risk register; capability-permission map (agents × tools × data tiers × purposes); knowledge-tier policy catalog as policy-as-code; MBOM register; delegation-chain audit map [Security Auditor §6]. New capability — capability-economics modelling (NPV per capability, optionality value, lock-in cost analysis, portfolio rebalancing) [Platform Economist §6]. New discipline — knowledge-stratum economics including the forgetting question.
- **Process (blue).** Continuous architecture review replaces or supplements quarterly ARB cycles for routine change; high-stakes decisions remain ARB-gated. Reference architectures shift from PDF prose to MCP-wired patterns. Capability registry becomes a board-visible asset. TOGAF 11 (probable 2029) and ArchiMate 4.0 (probable 2030) likely add agent and knowledge-tier semantics [Trend Analyst §1.5 BET].
- **Emotion (red).** Identity bifurcation is real. Classical EA self-concept ("I hold the model; I gate the change") and EA-3.0-paradigm self-concept ("I cultivate the orchestration substrate; I shape the routing policy") are different careers with different self-concepts, not a smooth evolution [HR Advisor §6]. The practitioner in 2026–2028 has to choose, and the choice is identity-deep. The vocabulary lags the reality — that lag is itself a source of professional anxiety the deliverable should name.

Old EA capability map vs. EA-3.0-paradigm capability map.

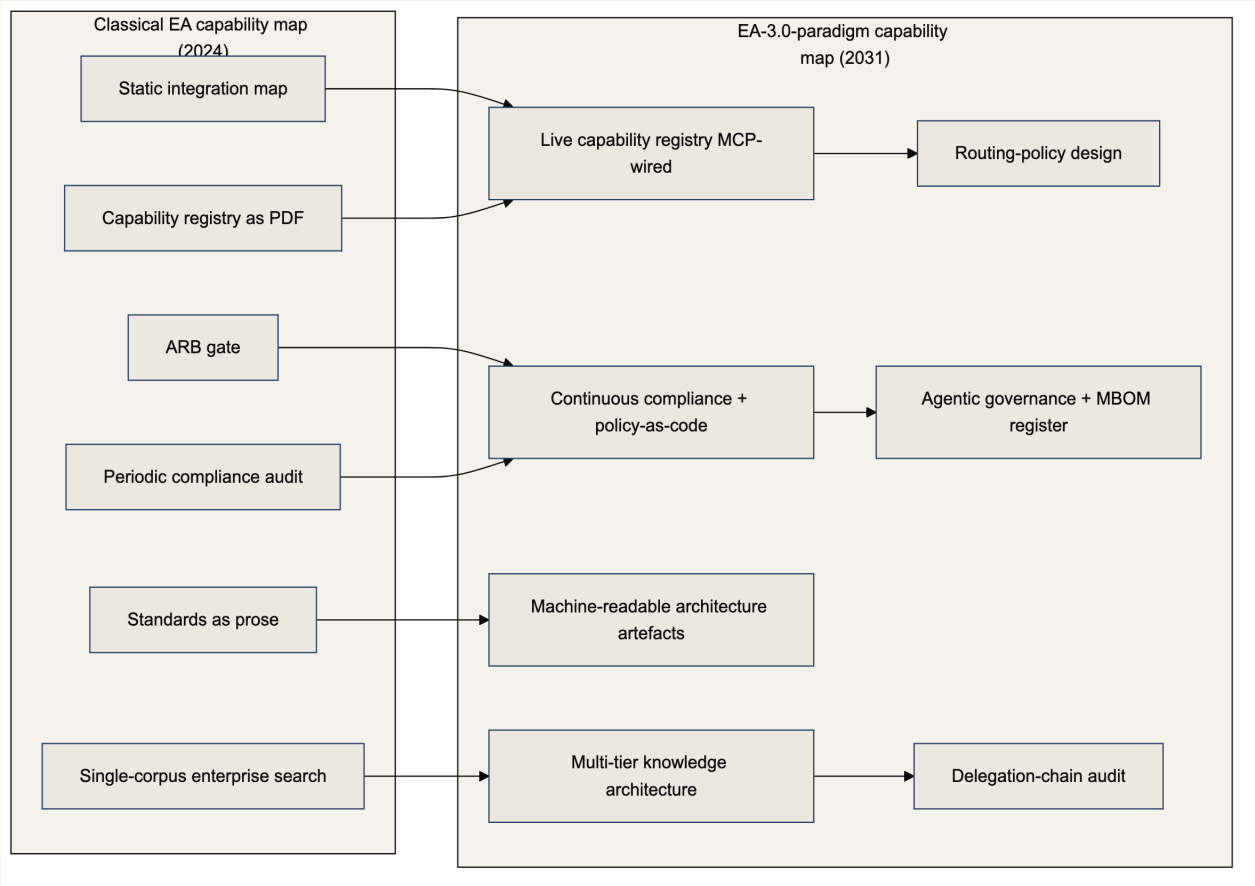

Tech participation. Assistant universal in EA tooling (Forrester EAMS Landscape Q4 2025 [T2]). Agentic core to the discipline by 2028. Autonomous bounded by regulatory and governance maturity. **Knowledge-tier role.** Governs topology, classification, governance gates. **Determinism-routing role.** Owns the architectural primitive; defines firm-level routing policy; mediates business and engineering [AI Architect §4.1].

4. The cognitive/identity shift through-line

The Copernican turn the foundational question names ("ajatuksen muutos") is the load-bearing through-line across all five populations. The shape repeats with high consistency [HR Advisor §1, §8]:

POPULATION	TOOL / OUTPUT IDENTITY (CLIFF)	DISCIPLINE IDENTITY (CONTINUITY)
P1 The firm	"We are a [vertical] company that uses technology"	"We are an outcome-delivering organism that orchestrates human + agent capability"
P2 Operational worker	"I'm a process executor with a Copilot"	"I am the judgment layer for the customer outcome"
P3 Leader-as-leader	"I'm a decision-maker who uses AI dashboards"	"I am the trust architect, routing-policy designer, sense-maker"
P4 Technical builder	"I'm a code author who uses Copilot"	"I am an engineer of system judgment; intent specifier; verifier"
P5 Whole-system architect	"I'm a TOGAF practitioner with AI EA tools"	"I am the architect of the orchestration substrate"

Reference-class data from prior platform shifts (mainframe → client-server → web → cloud → SaaS → lakehouse) is consistent: tool/output self-concepts face cliff-edged compression; discipline self-concepts cross the shift with craft intact and often elevated [baseline §D7, Trend Analyst §2, HR Advisor §8]. **The cliff is on tool / output identity. The continuity is on discipline identity** — the single sentence organising the deliverable.

Tool-identity → discipline-identity slope, by population, 2026 → 2031
0 = pure tool/output identity; 100 = pure discipline identity. Reference-class anchored.

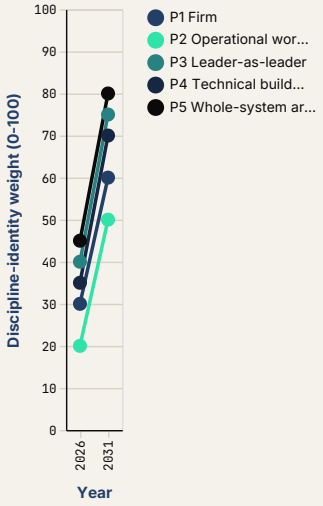

The 31-point gap between adoption and integration is an identity gap. Workers reach for the tool; they have not yet rewritten the story of their own value around it [HR Advisor §1]. Adoption is behavioural; integration is identity-level. Tooling investment without identity-language investment hits a ceiling. Successful 2026–2031 transitions seed the language for the new self-concept, not just the courses for the new tooling. **HR's role is to seed the vocabulary.** The lagging-edge firm in 2031 has the same Copilots, same Glean, same agentic workflows — but workers still describe themselves as task executors using a faster typewriter, leaders still as decision-makers using a smarter dashboard. The differentiator is the vocabulary the firm uses to describe what its people are for.

5a. Cross-cutting: Determinism routing

Determinism routing is the new EA meta-capability — the most important architectural decision the discipline acquires in this transition. The pattern is real and observable today [AI Architect §5, Security Auditor §1, §7].

Where determinism wins. Latency-critical surfaces (sub-100ms response budgets); cost-critical surfaces (per-transaction cost models where token cost is meaningful margin); audit-critical surfaces (financial close, regulatory filing); regulated decisions (Finnish administrative-law-bound under Chancellor of Justice OKV/2213/70/2024, ruling 2025-04-02; EU AI Act high-risk classifications; safety-critical control loops). **Where non-determinism wins.** High-variance contexts (customer support, exception handling, novel operational situations); language-heavy work (drafting, summarisation, translation, classification of unstructured input); discovery and synthesis; generative interfaces.

The bridge layer. Structured outputs and function calling; guardrails (Forrester AEGIS [T2], OPA Rego, AWS Cedar [T1, public standard]); agent frameworks with policy hooks (rate limits, cost budgets, jurisdictional routing); RAG with citations; hybrid retrieval-then-rule (LLM interprets intent and selects rule, rule executes — common in Finnish and Swedish public-sector architectures because the legal-interpretation half is rule-bound).

```
quadrantChart
    title Work types on determinism axis × volatility axis
    x-axis "Deterministic" ⟶ "Non-deterministic"
    y-axis "Stable workload" ⟶ "Volatile workload"
    quadrant-1 Agentic with guardrails
    quadrant-2 Agentic-first
    quadrant-3 Deterministic SOP
    quadrant-4 Deterministic with audit
    "Financial close": [0.1, 0.3]
    "Regulated decision Kela-class": [0.15, 0.25]
    "Customer support routine": [0.4, 0.55]
    "Exception handling": [0.7, 0.7]
    "Drafting and summarisation": [0.8, 0.4]
    "Generative interfaces": [0.85, 0.6]
    "Discovery and synthesis": [0.9, 0.85]
    "Dynamic pricing within bands": [0.55, 0.75]
    "Capacity scaling within SLA": [0.45, 0.8]
    "Compliance reporting": [0.2, 0.4]
```

Continuous compliance binds the routing decision. Policy-as-code (OPA Rego, AWS Cedar [T1, public standard]) expresses access, capability, retrieval, and output policies as version-controlled artefacts evaluated at runtime. Runtime enforcement at four points: agent invocation, capability call, retrieval, generation. Quarterly ARB cycles replaced by policy-conformance dashboards for routine change; high-stakes ARB persists. SOC 2, ISO 27001, FedRAMP become continuously attestable rather than periodic-audit-only [Security Auditor §7]. Regulatory wave under Regulation (EU) 2024/1689: prohibited-AI from February 2025; GPAI model rules from August 2025; high-risk Annex III from August 2026; full applicability from August 2027 [T1, public standard; CoVe Claim 10]. DORA and Finnish administrative-law push EA toward continuous compliance rather than periodic audit.

Strategic implication: **routing-policy design becomes a board-reportable competency by 2031** in lead-edge firms [Platform Economist §9 VISION; likely for leading-edge, speculative for median].

5b. Cross-cutting: Knowledge stratification

The hidden EA shift. Single-corpus enterprise search of 2024 is gone in 2031 in lead-edge firms; in its place is a tier-aware retrieval architecture with classification, governance gates, decay policies, and provenance baked into the runtime [AI Architect §3, Security Auditor §3, Platform Economist §2.6]. This is **not in TOGAF 10 today**.

Tier topology. Personal (employee device + memory) → Team (workspace, channel, project) → Function (HR, Finance, R&D) → Corporate (semantic layer + corporate KG) → Cross-org (partners, suppliers, regulators) → Public (web, published outputs). Each tier has its own retrieval permissions, contribution gates, governance review, and decay policy. The agentic memory layer (Letta, Mem0, Zep, Cognee primitives [T1]) lets an agent retain context across tiers and sessions; without it, each query is a cold start.

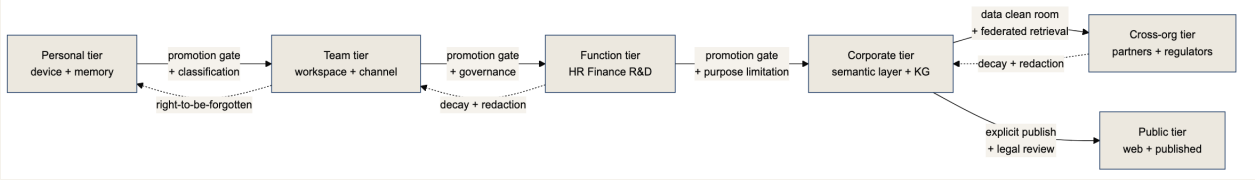

Required architectural primitives. Agentic memory (episodic, semantic, procedural). Persistent context (long-context models help but do not replace memory architecture). MCP-style addressability per tier. Federated retrieval. Knowledge provenance (source, timestamp, authoring identity, classification, confidence on every fact). Decay and forgetting policies (TTLs, right-to-be-forgotten propagation, supersession). Governance gates between tiers — where the **knowledge curator** role lives [Forrester T2].

Tier-boundary security. Every retrieval call carries a verifiable claim set (identity, tier membership, purpose, lawful basis); the retrieval engine evaluates policy before returning embeddings; output-time redaction as backstop; per-tier audit logs [Security Auditor §3]. Hard cases: personal-tier on personal devices (GDPR data-subject rights apply; employer access bounded by employment-law lawful basis and proportionality); team-tier inside multi-tenant platforms (open governance question); corporate-tier with personal data (GDPR Article 22 plus EU AI Act high-risk classification); cross-org (data-clean-room patterns, federated retrieval; no open standard in 2026).

Knowledge-stratum economics — the strategic implication. Corporate-tier knowledge under governance is **the sustainable moat** in an agentic world. Pricing playbooks, engineering tribal knowledge, customer-segment behavioural data, compliance interpretations are durable when structured, addressable, well-governed [Platform Economist §2.6]. Leaked through ungoverned agent calls, ungoverned fine-tuning, or third-party retention defaults, the same knowledge becomes commodity. The forgetting question is genuinely new — agent memory

persists in places EA has never had to govern (Letta state, Mem0 vector stores, Zep temporal graphs, third-party caches). Firms that get this wrong have knowledge-leakage incidents 2027–2029 that reshape their posture. By 2030, knowledge-stratum governance is a default RFP requirement [Platform Economist BET-4].

The pattern share across enterprise knowledge workloads shifts decisively over the horizon: vector-only RAG declines from ~62% (2026) toward ~18% (2031); GraphRAG and hybrid grow then plateau at ~38%; multi-tier agentic memory grows from ~6% to ~42% as the dominant pattern by 2031 in the speculative scenario [AI Architect §3.4 BET]. Adoption curve assumes compounding-capability trajectory. Under capability plateau, multi-tier memory share caps at ~20–25% by 2031 and vector-only RAG holds at 35–45% [bias-counter §4]. Shares overlap (a workload can use vector-only inside an agentic-memory wrapper); the chart shows primary pattern.

Knowledge-stratum pattern share across enterprise workloads, 2026 → 2031
Compounding-capability scenario. Under plateau: vector-only holds at 35–45%, multi-tier caps at 20–25%.

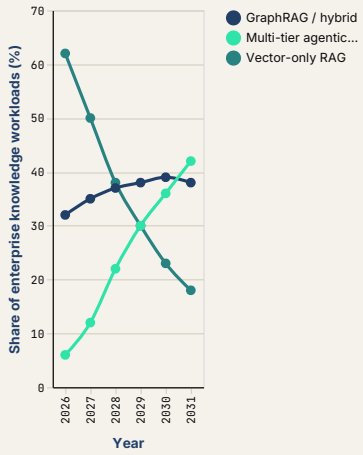

6. Reference-class — base rates from prior platform shifts

Six prior platform shifts inform the 2026–2031 forecast: mainframe to client-server; client-server to web; on-prem to cloud; manual QA to automated SDET; pre-DevOps to DevOps; and pre-data-engineer to the data-engineer role. Three operational shifts add complementary signal: bank teller post-ATM (Bessen [T2]), call-centre evolution, and the travel-agent collapse from roughly 140k US practitioners in 1990 to 70–80k by 2010 [T2]. Across all of them the pattern is consistent enough to treat as a base rate rather than an analogy.

The 5–10-year mainstreaming arc. Roughly five years separates a technology inflection from emerging-role visibility, and roughly ten years separates inflection from full mainstreaming with role-specific bootcamps, university programmes and clear career ladders [Trend Analyst §2; baseline §D7]. Cloud security took 8–10 years; DevSecOps the same; API security 4–6 (faster, narrower scope under regulatory pressure); zero-trust 4–6 (regulator push). AI-system security maps onto this curve at year 2–3 of an 8–10-year arc — meaning the 2026–2031 horizon **ends roughly when mainstream begins** for the median firm [certain on shape; likely on duration band; Security Auditor §8].

Old roles narrow rather than vanish. Mainframe operators, manual QA, on-prem sysadmins still exist in 2026 — narrowed, regulated, premium-priced for late-career holders. COBOL is the cautionary tail-end story: 84,000 unfilled positions by 2020 because of cliff-shaped university pipeline; average COBOL programmer age around 58 by the late 2010s [T2]. The travel-agent collapse is the cautionary case for roles whose value was throughput on codifiable tasks and that did not re-anchor on judgment, relationship, or trust [Operations Manager §6].

Mid-career holders without portable abstractions are most exposed. Senior practitioners gain a scarcity premium; junior practitioners reskill into the new role; the mid-career cohort that skilled up the old paradigm but cannot port abstractions is the most exposed slice [certain on shape; speculative on percentages; Trend Analyst §2]. The pattern compresses to one sentence: **tool/output identity compresses; discipline identity elevates.** It applies at the firm level too — the cloud reference class shows roughly 10–15% of firms reach genuinely cloud-native operating models even after 15 years; most reach lift-and-shift and stop. The agentic shift will likely follow the same shape — most firms reach Copilot-augmented application stacks but never reach capability-portfolio operating models [likely on shape; speculative on cohort size; baseline §D7].

Survival distribution. Pattern-anchored estimates from the seven prior shifts (caveat: pattern-anchored, not BLS-anchored — the BLS time-series back-extraction was out of scope for this engagement, see §12 unknowns): roughly **25–35% stay-in-role / 40–60% migrate / 15–30% exit** [likely on bands; speculative on point estimates; Trend Analyst §2]. New roles take 5–10 years to mainstream — the data-engineer arc went from emerging title in roughly 2015 to

standard role with role-specific bootcamps and university programmes by roughly 2025 [T2/T3]. The same arc applies to "agent integrator," "verifier engineer," "knowledge architect," "agentic governance champion" — emerging now, mainstream around 2030–2032.

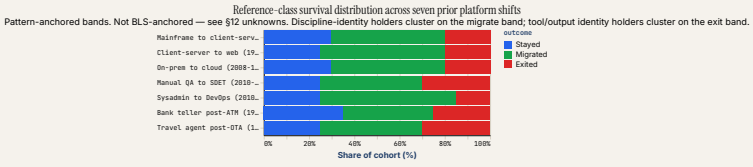

The chart is the empirical foundation under the cognitive-turn through-line: the migrate-and-stay cohorts cluster on discipline-identity holders; the exit cohort clusters on tool-identity holders who could not re-anchor. The 2026–2031 forecast for enterprise architecture and the four other populations applies these bands as base rates, not as point predictions.

7. Watch-list — falsifiable leading indicators

The watch-list is the operational falsification instrument. Each indicator is specific, dated, and grounded on a verifiable signal source — analyst publication, BLS taxonomy, vendor public registry, or named regulatory event. Trend Analyst §6 supplies the foundational ten; the synthesis adds three strategic-level indicators.

Foundational ten [Trend Analyst §6].

1. By 2027-Q2, **≥3 of AWS / Azure / GCP / Oracle publish first-party MCP server registries**, each with ≥50 servers and ≥3 first-party-published servers. Falsifies Scenario B if true; supports Scenario A or C.
2. By 2028-Q4, **Forrester Wave or Gartner Magic Quadrant for Enterprise Architecture reports ≥30% of EA roles include "AI system architecture" as top-3 responsibility** (currently <10% based on the 2024 Open Group Architecture Roles and Skills Series Guide). Supports Scenario A or C.
3. By 2029, **≥1 named Fortune 100 firm publicly restructures its EA function as an "agentic operating model"** or equivalent (named press release, not vendor-marketing claim). Strong supporter of Scenario C.
4. By 2031, **BLS computer systems analyst employment in the 480k–540k band** — close to the 2024 baseline of 521,100 with the +9% projection partly offset by AI substitution [T1]. Below 470k supports compression; flat 521k supports plateau; the band itself with internal task redistribution supports bifurcation.
5. By 2030, **≥3 distinct EA-adjacent job titles appear in LinkedIn Talent Insights and indeed.com taxonomies with ≥1,000 active postings each globally**. Direct support for Scenario C and Forrester's four-roles thesis.
6. By 2029, **AI-verkstad has ≥2 named referenceable production deployments cited in EU AI strategy reports or OECD Going Digital tracker** [T1, regeringen.se + skatteverket.se]. Independent confirmation of the cooperative-platform Nordic track.
7. By 2028-Q2, **an analyst report (Forrester / Gartner) documents ≥40% of Fortune 500 firms with ≥1 MCP-mediated agent in production workload** (not pilot or POC). Strong Scenario A indicator. The current T3 figure (78% from digitalapplied.com) is directional only and does not satisfy the indicator.
8. By 2030, **≥1 Nordic public-sector organisation is formally sanctioned for AI-system non-compliance** (Chancellor of Justice ruling or court judgment), becoming a region-wide reference case. Scenario-independent; structurally important for Nordic narrative.
9. By 2028-Q4, **Forrester EAMS Landscape names "agentic governance" and "knowledge stratification" as scoring criteria**. Scenario A or C supporter.
10. By 2028, **Stack Overflow Developer Survey shows AI-tool trust recovering above 70% (the 2023–2024 high), or falling below 55%** (further deterioration from the 2025 reading

of 60% [T2]). Direction is the indicator: recovery supports Scenario A; further fall supports Scenario B.

Strategic-level additions.

1. By 2027, **≥1 of McKinsey / BCG / Bain / Deloitte publishes a formal "agentic operating model" or "capability-economics" methodology** with named framework, executive curriculum, and case studies. Reference class: "platform strategy" 2015; "lean startup" 2010 [Platform Economist BET-2 + HR Advisor BET-2]. The absence of such a methodology by 2028 is the leading signal that the cognitive-turn timeline slips toward 2032–2034.
2. By 2028, **EU AI Act enforcement produces ≥1 high-profile fine ≥€100M against a Fortune 500 firm** [T1, Regulation (EU) 2024/1689]. The Act entered into force in August 2024. Prohibited-AI provisions applied from February 2025; GPAI model rules from August 2025; high-risk Annex III provisions from August 2026; full applicability across all provisions in August 2027. Reference class: GDPR enforcement 2019–2020 (Marriott, BA, Google) anchored market behaviour within roughly two years of full applicability; the equivalent anchoring window for the AI Act runs 2027–2029 [Security Auditor BET].
3. By 2029, **"enterprise capability registry / agent control plane" appears as a distinct procurement category in Forrester EAMS Landscape and Gartner Magic Quadrant**, displacing classical iPaaS / API management as the primary integration spend category at the modal mid-large enterprise [Platform Economist BET-1].

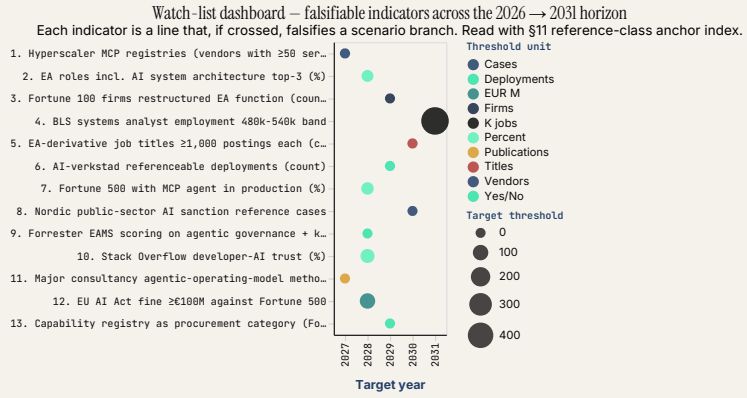

The watch-list is operational by design. A 2027 mid-horizon scenario re-grade evaluates indicators 1, 6 and 11; the 2028 audit evaluates 2, 7, 9, 10, 12; the 2029–2030 audit evaluates 3, 5, 6, 8, 13; the 2031 audit settles 4. Indicators that cannot be measured because no analyst publishes the relevant taxonomy, or because the BLS series is not yet released, surface in §12 as effort-gated unknowns. The watch-list operationalisation matrix — who measures, what cadence, what the publication lag is — is a follow-on artefact flagged for the 2027 mid-horizon review.

8. Capability-trajectory dependency

Every forecast in this document carries an implicit assumption about frontier-model capability through 2031. The honest answer in 2026 is that the trajectory is `[unknown]` and that this matters for some claims more than others. The dependency note below distinguishes the two camps explicitly so the reader can stress-test conclusions against either branch.

Robust to a 2027–2028 capability plateau — the foundation does not move if frontier capability flatlines. Protocol consolidation (MCP, A2A, AG-UI, A2UI under the Agentic AI Foundation [T1]) is a coordination problem, not a capability problem; the standards ship regardless of whether a frontier successor model arrives. Determinism routing as architectural primitive gains importance under plateau, because deterministic envelopes expand whenever agents are unreliable; hybrid retrieval-then-rule patterns common in Finnish and Swedish public-sector architectures become the dominant agentic pattern across regulated industries. The cognitive-turn through-line (tool/output identity compresses; discipline identity elevates) is reference-class-anchored across five prior shifts and does not depend on a particular AI capability ceiling — it slows but it still happens. Continuous-compliance regulatory pressure (DORA, EU AI Act, NIS2, Finnish administrative-law) arrives on regulatory clocks, not capability clocks. The Nordic public-sector cooperative-platform overlay is structurally distinct because of governance tradition, not capability.

Sensitive to a 2027–2028 capability plateau — these claims slip 1.5× to 2× later, or do not arrive by 2031 at all. Multi-tier agentic memory share at 42% by 2031 [AI Architect §3.4 BET] caps at roughly 20–25% under plateau; vector-only RAG holds at 35–45% rather than declining to 18%. Operational-worker compression magnitude at 50–65% over five years is the central case (band 45–60%); under plateau it lands in the 30–45% band and the second wave of autonomous-tier compression in 2030–2031 does not arrive. Leader cognitive-turn timing slips from a 2028–2029 inflection to 2032–2034. Frontier-firm compounding advantage narrows because the capability-access dimension collapses (mid-tier consumption converges with 2026 frontier on most enterprise tasks). Capability economics as a recognised C-suite topic [Platform Economist BET-2] does not materialise by 2031.

How the deliverable handles the dependency. Claims in §9 (Vision) and §10 (Bets) that depend on the compounding-capability scenario carry an explicit "compounding-capability dependency" tag. Probability bands per §3 (Scenarios A 25–35% / B 25–35% / C 40–50%) are the readable single number — they say the assumption has roughly two-thirds support, not certainty. The 2027 mid-horizon scenario re-grade is the honest checkpoint: by then the frontier-model price-per-token trajectory and the agentic long-horizon benchmark progress will discriminate compounding from plateau within reasonable precision [Dissent 2 in §12]. Until then the synthesis stays cohort-stratified rather than averaged.

9. Vision — what 2031 looks like

[VISION — argued, not sourced. Calibrated-confidence labels still apply. Capability-trajectory dependency note: this section assumes the central case (capability compounds at moderated pace through 2031); sensitive claims marked. Under the plateau branch, timelines slip 1.5–2× and several specific claims do not arrive by 2031.]

In 2031 the picture is a coherent bifurcation [Scenario C, working assumption with band 40–50%], with a Nordic public-sector overlay running as a third structurally distinct track.

Frontier — 15–25% of large enterprises [VISION; speculative on cohort size; likely on direction]. Boards review a capability portfolio the way they review a product portfolio today, with NPV-per-capability, lifecycle stage, governance tier, external-exposure status. CFOs allocate to capability domains, not departments; CTOs and Chief Capability Officers run the capability registry the way today's CDOs run the data lake — a contested, governed, board-visible asset. Integration is MCP-first; greenfield enterprise systems expose MCP servers as their primary machine-consumed interface, with OpenAPI persisting as the human-contract layer. Knowledge is stratified by default — multi-tier retrieval with classification, governance gates, decay, and provenance baked into the runtime; vendor convergence has consolidated to two or three enterprise-grade platforms; knowledge architecture is a recognised EA sub-discipline alongside data, information, and security architecture. Determinism is a routing decision, not an assumption — the firm publishes its routing policy and contests it in board meetings.

The CEO at a frontier firm describes herself as a trust architect and routing-policy designer — not a decision-maker — and the board accepts this language [HR Advisor §2; likely for leading-edge cohort by 2031; speculative for median Fortune 500; compounding-capability dependency]. The chief operating officer talks about agent oversight as a strategic competency on par with people leadership. Mid-level managers (newly named "agent orchestrators") manage hybrid teams in which the headcount measure is partly human FTE and partly agent-capacity-units. Operational workers describe their jobs in outcome terms: "I make sure the customer's first-week experience holds." Senior engineers describe themselves as judgment-of-systems specialists and the apprenticeship pipeline for junior engineers has been re-invented through curated learning sandboxes [HR Advisor BET-3, speculative]. Enterprise architects have bifurcated into specialist tracks (value mappers, digital twin strategists, knowledge curators, agentic governance champions) plus an "ecosystem cultivator" generalist track plus a residual classical-TOGAF cohort in regulated environments [Forrester taxonomy + synthesis].

Every consequential agent is in an agent inventory with purpose, capability scope, data tiers, and EU AI Act risk classification; every production agent has a model bill of materials, signed artefacts, and pinned versions; SOC 2 / ISO 27001 / EU AI Act conformity is attested from the

running system rather than reconstructed quarterly. The frontier firm has experienced at least one agentic-security incident with operational consequences and runs incident-response playbooks specific to prompt-injection, memory-poisoning, and delegation-chain compromise [Security Auditor §11; incident-window probability re-calibrated to 40–55% per bias-counters §5; compounding-capability dependency on incident timing, not on incident shape].

Modal Fortune 500 — the other 75–85% [VISION; likely on shape; speculative on cohort size]. Same Copilots, same Glean, same agentic workflows — structurally similar in tooling, but the identity language has not shifted. TOGAF 11 (released 2029) and ArchiMate 4.0 (released 2030) are the EA framework foundation; agent and knowledge-tier semantics have been added but the modal practitioner still describes themselves in classical-EA terms. Integration depth has risen from 44% to perhaps 60% — the deep-integration ceiling holds. The trust ceiling holds at 60–65% [T2, Stack Overflow trajectory] — agent suggestions are accepted at a higher rate but verification overhead is a permanent feature of the workflow. The agent boss pattern is mainstream for operational workers; service-ops headcount is 25–35% below 2025 levels in Fortune 500 firms. Workers still describe themselves as task executors using a faster typewriter; leaders still describe themselves as decision-makers using a smarter dashboard. These firms experience the agentic shift as cost compression without organisational learning, headcount cuts without identity reframing, and a measurable engagement and retention deficit.

Nordic public-sector overlay — third structurally distinct track [VISION; speculative on EU-reference outcome; likely on national-level shape]. Sweden's AI-verkstad (Försäkringskassan + Skatteverket, 200M SEK over 2026–2030 [T1, regeringen.se + skatteverket.se]) has become a referenceable global pattern — formally referenced in EU-level AI infrastructure documents [BET, speculative; reference class GAIA-X]. Finland's Chancellor of Justice ruling on Kela (OKV/2213/70/2024, 2025-04-02 [T1, finlex.fi + Yle]) has hardened public-sector AI procurement toward continuous compliance and explainability. The Nordic pattern is structurally distinct from US/UK — cooperative-platform model rather than vendor cloud-marketplace adoption; slower compression, larger residual SOP niche, union and public-sector protections shaping the rate. At least one Nordic government employer has published a "civil-servant self-concept" framework explicitly addressing the identity question for public-sector knowledge workers [HR Advisor §10 BET-5, speculative].

The defining feature of 2031 is not that AI is everywhere — it is that the integration-depth gap and the trust ceiling are now permanent operating-model features, not transitional discomforts. That sentence holds across both the central case and the plateau branch; it is the load-bearing claim of the VISION.

10. Bets — non-obvious calls

[BETS — explicitly speculative. Falsifiers attached. Calibrated-confidence labels. Compounding-capability dependency tagged where applicable.]

Eight calibrated bets. Each is concrete, dated, and falsifiable.

1. **By 2029, MCP-style protocols are the dominant integration paradigm in greenfield enterprise systems for frontier-firm new builds** [BET, likely for greenfield software-vendor systems; speculative for greenfield Fortune 500 IT; AI Architect §8 BET-1 with bias-counters §5 amendment]. The 13-month vendor-neutralisation arc (Anthropic Nov 2024 → AAIF Dec 2025; Microsoft Copilot Studio adoption Jul 2025 with 45M cumulative SDK downloads at that point; AWS adoption Nov 2025 with 68M [T1, T2]) is the strongest single anchor. **Falsifier:** if MCP adoption stalls below 30% of greenfield Fortune 500 enterprise APIs in the Forrester / Gartner annual integration survey by 2029-Q4, the bet fails.
2. **By 2030, ≥3 distinct EA-derivative job titles will have ≥1,000 active LinkedIn postings each globally** [BET, likely; Trend Analyst BET-1]. Concretely: "agentic operating model architect," "knowledge architect," and "AI governance architect." Reference class: data-engineer's decade arc. **Falsifier:** if no major job-board taxonomy carves out three EA-derivative titles by 2030 — or if any of the three has fewer than 500 postings — the bet fails.
3. **By 2028, EU AI Act enforcement produces ≥1 high-profile fine ≥€100M against a Fortune 500 firm** [BET, likely; Security Auditor §10 BET]. Reference class: GDPR enforcement 2019–2020 (Marriott, BA, Google). The Act's full applicability landed in August 2027; a roughly 12–18-month enforcement-anchoring window puts the first major fine in the 2028–2029 corridor. **Falsifier:** no enforcement action of that magnitude by end-2028.
4. **By 2030, "capability economics" is a recognised C-suite topic, with ≥1 major management consultancy publishing a formal methodology and ≥1 major business school running a capability-economics elective** [BET, speculative; compounding-capability dependency; Platform Economist BET-2]. Reference class: "platform strategy" 2015. **Falsifier:** no published methodology and no business-school course by end-2030.
5. **Bifurcation produces a measurable productivity gap between Frontier and Fast-follower firms of ≥1.5× in scaled-AI sectors by 2030** [BET, speculative; compounding-capability dependency; Strategy Consultant addition]. Productivity measured as AI-attributed revenue-per-employee or EBITDA-margin contribution in segments where AI is the primary lever (financial services, software, media). Reference class: cloud-native frontier-firm gap circa 2018 (Capital One, Netflix, Spotify, John Deere). **Falsifier:** if, by 2030, no Forrester / Gartner / McKinsey study identifies a frontier-vs-mainstream productivity gap of ≥1.5× in any scaled-AI sector, the bet fails — and Scenario C narrative weakens toward Scenario B.

6. **By 2029, the median Fortune 500 firm has rewritten career architecture for ≥1 functional area to recognise "agent orchestration" as a senior-track competency** — formally, in titles, level definitions, pay bands [BET, likely on shape; speculative on precise threshold (45–60% range plausible); HR Advisor §10]. Leading-edge cohort by 2027–2028; median by 2029. **Falsifier:** if Mercer / Korn Ferry / Aon levelling-survey data shows no formal agent-orchestration competency in any Fortune 500 functional career architecture by 2030, the bet fails.
7. **By 2030, runtime-attestable continuous compliance replaces point-in-time SOC 2 as the dominant assurance model for agentic-system vendors, with Big Four restructuring practice areas** [BET, speculative; compounding-capability dependency on agentic-deployment pace; Security Auditor §10 BET]. **Falsifier:** SOC 2 Type II remains the default vendor assurance artefact in Fortune 500 procurement RFPs by end-2030.
8. **By 2031, the Nordic public-sector cooperative-platform pattern (AI-verkstad style) is formally referenced as a model in ≥1 EU-level AI infrastructure document** (Commission communication, Council recommendation, AI Office working paper) [BET, speculative; Trend Analyst §9 BET-3]. Reference class: GAIA-X — also Nordic-and-Continental, also cooperative — which had partial uptake. AI-verkstad's combination of administrative-law anchoring, joint-government commitment (200M SEK over 2026–2030), and Nordic public-sector EA tradition makes it a stronger reference candidate. **Falsifier:** no formal EU-level reference by end-2031.

---

11. Reference-class anchor index

Every forward claim in §9 and §10 is mapped here to its reference-class anchor. Claims without anchors carry an additional `[speculative]` flag because reasoning forward without a base rate is the structural failure mode the discipline guards against. Anchor strength is **strong** (multiple prior shifts, well-documented), **weak** (one analogous shift or partial coverage), or **absent** (no reference class — pure projection).

FORWARD CLAIM	REFERENCE-CLASS ANCHOR	ANCHOR STRENGTH
Protocol consolidation (MCP / A2A / AG-UI / A2UI) by 2027–2028	Kubernetes-to-CNCF arc roughly 36 months, MCP-to-AAIF 13 months [T1, T2]	strong
EA discipline survives in transformed form	Five prior platform shifts (mainframe, client-server, web, cloud, SaaS, lakehouse) — every one produced bifurcation, none produced dissolution	strong
Operational-worker compression 45–60% band over 5 years	Bank teller post-ATM (Bessen [T2]), call-centre evolution, manufacturing post-robotics, travel-agent collapse	strong
Mid-career most exposed; senior premium, junior reskill	Reference-class evidence across all seven prior shifts [Trend Analyst §2]	strong
AI-system security mainstream 2031–2033	Cloud security 8–10y, DevSecOps 8–10y, API security 4–6y, zero-trust 4–6y	strong
Old roles narrow rather than vanish	COBOL pattern, mainframe operators, manual QA, on-prem sysadmins	strong
New role mainstream 5–10 years after emergence	Data engineer 2015→2025 arc [T2/T3]	strong
EU AI Act enforcement anchoring by 2028	GDPR enforcement 2019–2020 (Marriott, BA, Google) within 24 months of full applicability	strong

FORWARD CLAIM	REFERENCE-CLASS ANCHOR	ANCHOR STRENGTH
MBOM mainstream by 2029	SBOM after EO 14028 (US Executive Order on Cybersecurity, 2021 [T1, public standard])	strong
Capability economics as C-suite topic by 2030	"Platform strategy" 2015 (Eisenmann, Parker, Van Alstyne)	weak
Three-mode firm posture (Frontier / Fast-follower / Specialist)	Cloud-era frontier-vs-mainstream pattern 2010–2018 [Trend Analyst §5]	strong
Frontier 15–25% / Modal 75–85% specific cohort sizing	Cloud-native frontier-firm cohort (Capital One, Netflix, Spotify, John Deere) — pattern shape only	weak
Bifurcation (two-speed enterprise) as modal trajectory	Cloud bifurcation 2010–2018 — partial fit; agentic shift may compress or amplify	weak
Multi-tier agentic memory 42% share by 2031	None — pure forward projection on AI Architect §3.4 BET	absent
Knowledge stratification as recognised EA sub-discipline	Data architecture's emergence 1990s–2000s — partial fit	weak
Nordic cooperative-platform pattern as third track	GAIA-X (cooperative European cloud federation, partial uptake 2019–2024)	weak
Leader cognitive turn 2028–2029 (median 2030–2032)	Digital-transformation leader identity 2014→2018–2020 leading-edge, still incomplete at median 2024	weak
Routing-policy design as board-reportable competency	Cybersecurity literacy as board competency post-Equifax 2017 — partial fit	weak
Verifier-engineer specialism consolidating by 2030	Data-engineer arc — same shape, different domain	weak

FORWARD CLAIM	REFERENCE-CLASS ANCHOR	ANCHOR STRENGTH
Apprenticeship pipeline re-invented in lead-edge firms by 2031	Post-CI/CD apprenticeship structures, post-DevOps SRE training — partial fit	weak
Frontier-vs-Fast-follower productivity gap ≥1.5× by 2030	Cloud-native productivity-gap data 2015–2020 — directionally supportive, exact multiplier unsourced	weak
Civil-servant self-concept framework (Nordic) by 2031	None — structural argument from administrative-law tradition	absent

The pattern is clear and useful for the reader. **Strong-anchored claims** (eleven of twenty-two) are the load-bearing claims of the deliverable; they survive both the compounding and plateau branches per §8. **Weak-anchored claims** (nine of twenty-two) carry their `[speculative]` or `[BET]` labels honestly — the shape is supported but the sizing is not. **Absent-anchor claims** (two of twenty-two) are flagged additionally because they extrapolate without a base rate; the reader should treat them with the heaviest scepticism. Multi-tier agentic memory at 42% by 2031 is the most consequential absent-anchor claim — it carries a compounding-capability dependency note and caps at 20–25% under plateau per §8.

12. Dissent register and unknowns

Dissent is preserved, not averaged. Specialists disagreed on ten substantive points; the five most consequential are summarised here. The full register accompanies this document as a companion artefact.

Top five dissents.

1. Vendor concentration vs open-protocol fragmentation 2027–2029 [AI Architect vs Platform Economist]. Either MCP / A2A / AG-UI under Linux Foundation governance hold the protocol layer open, or hyperscaler control planes (Copilot Studio + M365 + Azure AI Foundry; Bedrock + Q; Agentspace + Gemini Enterprise) capture the layer above and lock-in moves up the stack. **Held as:** Cross-exam Pair 3 layered resolution — protocol open, lock-in concentrates one layer up; both positions remain alive on whether protocol substrate holds long enough. `[unknown]`.
2. Foundation-model cost curve through 2031 [optimistic compounding vs plateau]. Either frontier-model inference cost falls roughly 10× per two years on automatable cognitive load through 2031, or it flattens 2027–2028 (compute-bound, energy-bound, regulation-bound). **Held as:** robust-vs-sensitive split per §8; capability-trajectory dependency note attached to all sensitive claims. `[unknown]`.
3. Autonomous-tier deployment scaling 2029–2031 [optimistic vs cautious]. Either bounded autonomous decisions become standard for routine-heavy operational segments by 2029–2031, or the 2027–2029 incident window resets industry caution and continuous-compliance maturity lags by 3–4 years. **Held as:** cohort-stratified — Frontier 15–25% reach mature autonomous tier in routine-heavy segments; mainstream Fortune 500 stays in agent-boss + assisted-autonomous mode pending compliance maturity. `[unknown]`.
4. Mid-career operational-worker retraining success rate [HR Advisor / Operations Manager optimistic vs pre-mortem on P2]. Either 30–50% navigate to elevated self-concept and 20–35% exit, or compression is 75%+ and judgment-elevation capture is 10–20%. **Held as:** `[unknown — no clean reference-class data]`; carry the 30–50% / 20–35% / residual band but flag explicitly as the optimistic-management midpoint; bands tightened to 45–60% per bias-counters §5.
5. EA discipline survival vs dissolution [majority survival-in-transformed-form vs Path D dissolution]. Either Forrester's four emerging EA roles plus an ecosystem cultivator track plus a residual classical cohort, or the Enterprise Architect title narrows sharply by 2031, concentrated in regulated public-sector and legacy stewardship niches. **Held as:** Cross-exam Pair 1 resolution stands `[likely → certain]` on survival, `[speculative]` on four-paths sizing; minority dissolution outcome remains visible — averaging it out is exactly the failure mode the dissent register guards against.

Top five unknowns.

1. Foundation-model economics through 2031. No anchored 2026–2031 forecast for frontier-model inference cost-per-standardised-task. Public benchmark progress saturates within 12–18 months of release. `[Inherently speculative — forecasting limit]`.
2. Integration-depth trajectory. The 84% adoption / 44% integration / 33% scaled / 60% trust gap is one snapshot. Whether it closes (Scenario A), persists (Scenario B), or bifurcates (Scenario C) is the central scenario-discriminating question. `[Resolvable when 2027-2028 successor surveys land — Stack Overflow / JetBrains / Microsoft WTI / McKinsey State of AI]`.
3. Agent reliability in production. No consistent T1/T2 benchmark for override rate, hallucination-detection rate, cost-per-task in regulated workloads. `[Resolvable when data exists — vendor reporting from ServiceNow / Salesforce / Microsoft customer-ops platforms; survey question in WTI 2027-2028]`.
4. Reference-class quantitative gaps. BLS time-series cohort-survival breakdowns at year +5 and +10 for prior shifts not centrally published. `[Resolvable with effort — primary BLS Current Population Survey microdata + occupational matrix back-extraction would close it]`.
5. Watch-list operationalisation. Indicator-by-indicator mapping of who measures, what cadence, what data product publishes it, what the publication lag is. `[Resolvable with effort — research task for the 2027 mid-horizon scenario re-grading]`.

The unknowns concentrate in three structural classes: time-gated (resolve as next-generation surveys / studies / regulatory actions arrive 2027–2030), effort-gated (primary research could close them now), and forecasting-limit (genuinely speculative because dependent on capability and economics trajectories that 2026 evidence cannot anchor). The deliverable treats all three honestly.

Glossary

MCP (Model Context Protocol). Open protocol for AI model interaction with tools and contextual data; Anthropic-released November 2024, donated to the Linux Foundation Agentic AI Foundation December 2025 [T1].

A2A (Agent-to-Agent). Protocol family for inter-agent communication and capability negotiation; companion to MCP under AAIF governance.

AG-UI / A2UI. Open protocols for agent-rendered user interfaces. AG-UI is CopilotKit-stewarded with about sixteen events over SSE / WebSocket; A2UI is Google's declarative generative-UI specification [T1].

RAG (Retrieval-Augmented Generation). Architectural pattern where an LLM grounds its output in retrieved documents rather than parametric memory alone. Vector-only RAG is the dominant 2024–2026 implementation.

GraphRAG. RAG variant that uses a knowledge graph plus community summarisation rather than vector retrieval alone; LazyGraphRAG (Microsoft Research, June 2025 [T1]) drops indexing cost to roughly 0.1% of full GraphRAG.

Agentic memory (episodic / semantic / procedural). Memory primitives that let an agent retain context across sessions and tiers. Episodic memory stores discrete events; semantic memory stores generalised facts; procedural memory stores how-to knowledge. Productised in Letta, Mem0, Zep [T1].

Knowledge stratification. Multi-tier knowledge architecture spanning personal, team, function, corporate, cross-org, and public tiers, each with its own retrieval permissions, contribution gates, governance review, and decay policy. Replaces the single-corpus enterprise search of 2024.

Determinism routing. The architectural decision about which surfaces stay deterministic (rule-based, audit-traceable) and which move to probabilistic agentic execution. Engagement-internal vocabulary; the underlying capability is real and observable today.

Three-threads framing. Synthesis structuring of the 2026–2031 transition as three loosely-coupled threads — agentic orchestration (protocols), knowledge stratification (data-to-memory), determinism routing (engineered choice) — maturing on different curves rather than as one paradigm. Replaces the engagement-internal "EA 3.0" label.

Agent boss (Microsoft framing). Term coined in Microsoft Work Trend Index 2025 [T1] for someone who builds, delegates to, and manages agents. Used as descriptive role today; speculative whether it consolidates as a formal job title by 2030 [Operations Manager BET-2].

Continuous compliance. Architectural discipline where policy-as-code (OPA Rego, AWS Cedar [T1, public standards]) is evaluated at runtime against agent invocations, capability calls, retrieval, and generation. Replaces periodic audit with continuously attestable conformance to SOC 2 / ISO 27001 / EU AI Act.

Capability registry. Enterprise inventory of agent-callable capabilities — purpose, owner, data tiers, governance class, EU AI Act risk classification. Becomes a board-visible asset in Frontier-mode firms by 2031 [VISION].

Adoption-vs-integration gap. The empirical gap between AI tool use (84–85% [T2]) and workflow integration (44% [T1]); between organisational AI use (88% [T2]) and scaled AI (33% [T2]); between 2023–2024 trust at 70% and 2025 trust at 60% [T2]. The single most important data point in the synthesis.

Bifurcation (two-speed enterprise). Scenario C from Trend Analyst §5; band probability 40–50% per bias-counters §1. Frontier firms compound integration depth while the median Fortune 500 stays in surface adoption; the cohort gap widens through 2031.

AI-verkstad (Sweden). Joint administrative AI infrastructure for Swedish public administration, run by Försäkringskassan and Skatteverket; 200M SEK total, 100M SEK each, 2026–2030 [T1, regeringen.se + skatteverket.se]. Reference candidate for an EU-level cooperative-platform pattern by 2031.

Source register summary

The deliverable rests on three tiers of evidence; the full source register accompanies this document as a companion artefact. T1 (primary sources) totals 47 entries — official documentation, peer-reviewed papers, regulatory filings, and primary government sources. T2 (major analyst, industry, quality journalism) totals 32 entries — Gartner, Forrester, McKinsey, Microsoft Work Trend Index, GitHub Octoverse, Stack Overflow Developer Survey 2025, JetBrains State of Developer Ecosystem 2025, peer-reviewed-adjacent publications. T3 (practitioner blogs, vendor self-reports, industry magazines) totals 22 entries — used directionally only, never as foundational evidence.

Key T1 anchors named. Microsoft Work Trend Index 2025 (75% gen AI use; "agent boss" coined). McKinsey State of AI 2025 (88% AI use; 33% scaled). JetBrains State of Developer Ecosystem 2025 (44% workflow integration). Stack Overflow Developer Survey 2025 (84% AI tool use; 60% trust). GitHub Octoverse 2024 + 2025 (80% of new developers use Copilot in week one). GitHub Copilot RCT (Peng et al., arXiv 2302.06590, 55.8% faster, 95% CI 21–89%). Anthropic MCP introduction + AAIF donation (Linux Foundation, December 2025). Microsoft Agent Framework GA Q1 2026 (AutoGen + Semantic Kernel merger October 2025). Glean ($200M ARR, $7.2B valuation, June 2025 Series F). Microsoft GraphRAG and LazyGraphRAG. Letta, Mem0, Zep (agentic memory primitives). Försäkringskassan + Skatteverket AI-verkstad (200M SEK, 2026–2030, regeringen.se + skatteverket.se). Chancellor of Justice OKV/2213/70/2024 ruling on Kela (2025-04-02, finlex.fi + Yle). EU AI Act, Regulation (EU) 2024/1689 (Feb 2025 prohibited / Aug 2025 GPAI / Aug 2026 high-risk Annex III / Aug 2027 full applicability). BLS Computer Systems Analysts OOH (521,100 in 2024; +9% projected). arXiv 2510.22003 (33-study GenAI's impact on EA in agile literature review). arXiv 2502.17443 (agentic workflows + enterprise APIs).

Source-tier discipline. Vendor self-reports (CrewAI 60%, Glean 93% adoption rate, "78% of enterprise AI teams report MCP-backed agent in production") are retained as T3 directional only and do not carry foundational claims. Vendor names introduced for sectoral colour (Salesforce Agentforce, ServiceNow Now Assist, Atlassian Rovo, Drata, Vanta, Secureframe, Langfuse, LangSmith, Arize, Helicone, Bizzdesign, LeanIX, Ardoq, MEGA) carry inline `[T2/T3, sectoral illustration; not separately verified]` labels. Public standards (OPA Rego, AWS Cedar, NIST 800-207, OWASP LLM Top 10, MITRE ATLAS, EO 14028) carry inline `[T1, public standard]` labels at first use. The deliverable's load-bearing claims rest on T1/T2 evidence; T3 is signal, not foundation.