

— THE AGENTIC ORGANIZATION — REPORT I

Building the *Organizational Agentic Layer*

The next era of enterprise software is not another app. It is a coherent layer of intelligent agents that sits across your business, mobilizes its context, and multiplies the output of the team you already have.

AUDIENCE
C-suite · mid-market & SMB

PUBLISHED
April 2026 · Helsinki

READING TIME
14 minutes

AUTHORS
Linkenite Research

Five conclusions from the 2025–2026 data

KEY FINDINGS

- | | | | | |
|---|---|---|--|---|
| <p>01</p> <p>Adoption is universal. Coherence is rare.</p> <p>88% of organizations use AI in at least one function. Fewer than 10% have scaled agents in any single function.</p> | <p>02</p> <p>The piecemeal era is expensive.</p> <p>Over 90% of employees use AI at work, yet only 40% of employers have an official subscription. Value leaks from every edge.</p> | <p>03</p> <p>The performance gap is widening fast.</p> <p>Future-built firms post 1.7× revenue growth, 1.6× EBIT margin, and 3.6× three-year TSR of laggards.</p> | <p>04</p> <p>The underlying stack is finally ready.</p> <p>Frontier agentic capability jumped from 42% to 77% on SWE-bench in four months. Inference prices fell 50× annually.</p> | <p>05</p> <p>Mid-market can lead, not follow.</p> <p>Mid-market firms ship production agents in 90 days. Enterprises need 9+ months. The window to compound that lead is now.</p> |
|---|---|---|--|---|

§ 01 — THE PROBLEM The piecemeal LLM problem

Most organizations are not under-using AI. They are *incoherently* using it. One person drafts emails in ChatGPT. Another lives in GitHub Copilot. A third pastes meeting transcripts into Gemini for summaries. A marketing contractor runs prompts through Claude. Customer support has a standalone chatbot that has never spoken to the CRM.

The result is a kind of cognitive archipelago — islands of intelligence with no connective tissue.

The data tells the story plainly. **Reco's 2025 State of Shadow AI Report** found that companies with 11–50 em-

ployees average 269 unsanctioned AI tools per 1,000 employees. Harmonic Security analyzed 22.4 million enterprise prompts in 2025 and concluded only 40% of companies run AI through official subscriptions. The rest is free-tier, personal-account, copy-paste.

This is not a compliance story. It is an operating-model story. **Every scattered interaction produces a unit of output with no memory, no audit trail, and no organizational context.** An SDR who asks ChatGPT to write a follow-up learns nothing the CRM does not already know. The CRM learns nothing about the decision the SDR just made. The next person researches the same account from scratch.

//

95% of enterprise generative AI pilots deliver no measurable P&L impact. The 5% that do share one trait — the AI is woven into an end-to-end workflow, not bolted onto a task.

MIT's NANDA initiative, in its July 2025 *State of AI in Business* report, captured the cost of fragmentation with one of the year's most-cited numbers. Among the 5% of pilots that do deliver P&L impact, a single factor predicts

success — whether the AI is woven into an end-to-end workflow rather than bolted onto a discrete task. The unlock is not better prompts. It is a coherent layer.

EXHIBIT A Adoption is universal. Scaled agentic use — and sanctioned access — are not.

ADOPTION

88%

of organizations use AI in at least one business function.

AT SCALE

≤10%

have scaled AI agents inside any individual function.

SHADOW AI

269

unsanctioned AI tools per 1,000 employees at 11–50 person firms.

Sources: McKinsey, *State of AI 2025*, November 2025; Reco, *State of Shadow AI*, 2025; Harmonic Security, 2025; BlackFog, January 2026.

SIDEBAR · SHADOW AI BY THE NUMBERS

- 90%+ of employees use AI at work; 40% of employers have an official subscription (Harmonic Security, 2025)
- 49% of workers use unsanctioned AI tools (BlackFog, January 2026)
- 46% of organizations reported data leaks through generative AI in 2025 (Cisco)

§ 02 — THE CONCEPT

What is the organizational agentic layer?

An **organizational agentic layer** is a coherent layer of software agents that sits across a business — watching its inboxes, calendars, meeting transcripts, documents, and systems of record — and that can read, reason, and write back into those systems on behalf of humans. It differs in kind, not degree, from what came before.

Piecemeal LLM use is a hammer in the hands of one person. Traditional workflow automation — the kind built on triggers and branching logic — is rigid plumbing that breaks when inputs do not fit its schema. Robotic process automation mimics mouse clicks on legacy screens. Each remains useful in narrow contexts. None carries *organizational context* across tasks, or adapts its behavior to the state of the world at the moment of action.

The agentic layer is different because its orchestrator — typically a terminal- or CLI-based runtime such as **Claude Code** or **OpenAI's Codex CLI** — gives a language model native access to tools, files, APIs, and other agents. It calls functions, queries databases, reads a Slack thread, cross-references a CRM record, drafts a reply, waits for a human approval, sends it, and logs the outcome. It does all of this inside a single, governable control plane.

Four layers make the pattern work

Layer 1 — Context. Agents ingest inbox, calendar, CRM, documents, meeting transcripts, ticketing systems, and

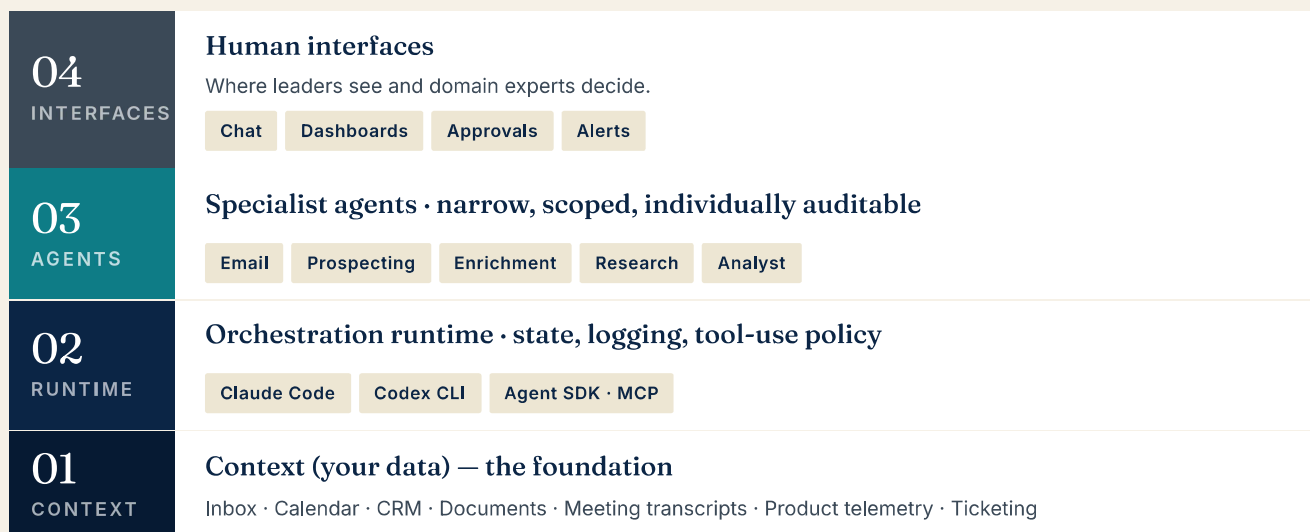
product telemetry. The **Model Context Protocol**, introduced by Anthropic in November 2024 and adopted by OpenAI in March 2025, is the emerging open standard. By its first anniversary MCP had surpassed 97 million monthly SDK downloads and 10,000 public servers, and was donated to the Linux Foundation's Agentic AI Foundation in December 2025.

Layer 2 — Orchestration. A terminal-native runtime coordinates the work — spawning subagents, holding state, calling tools, checkpointing progress. Claude Code reached roughly \$1 billion in annualized revenue within six months of general availability — evidence that CLI orchestration has become the dominant production pattern for non-trivial agent work.

Layer 3 — Specialist agents. Narrow, well-scoped agents handle discrete jobs: email triage, outbound sequencing, research, analytics, contract review. Anthropic's own multi-agent research system — an orchestrator that spawns parallel sub-researchers — outperformed a single-agent baseline by 90% on internal evaluations. The right pattern is *many small agents, not one god-agent*.

Layer 4 — Human interfaces. Chat, dashboards, and approval flows keep humans in the loop. The agent is the worker; the human is the editor-in-chief.

EXHIBIT B The organizational agentic layer sits between your data and your people.



Linkenite analysis. MCP is an open standard introduced by Anthropic (Nov 2024), adopted by OpenAI, Google and Microsoft (2025), donated to the Linux Foundation (Dec 2025).

§ 03 – WHY NOW Three curves crossed in 2025

THE CAPABILITY CURVE


Claude Sonnet 4, released May 2025, scored 42.2% on SWE-bench Verified. Claude Sonnet 4.5, released September 2025, scored 77.2%. Four months, a leap of roughly 75 percent of the original baseline on a widely respected agentic benchmark. **GPT-5**, August 2025, reached 74.9%. **Gemini 3 Pro**, November 2025, scored 76.2% and shipped with a one-million-token context window. Stanford's *AI Index 2026* reports that agent task

success on Terminal-Bench rose from 20% to 77.3% year over year. Long-horizon reliability was the gating problem for agents. It is now being solved in the field.

THE COST CURVE

Epoch AI's March 2025 analysis of inference prices at fixed performance levels found a **median 50x annual decline**, accelerating to **200x per year since January 2024**. What cost a dollar in inference in 2023 costs a cent today. What costs a cent today will likely be free-tier by 2027.

EXHIBIT C Capability rising, cost falling – in the same twelve months.



CAPABILITY

42 → 77%

SWE-bench Verified, four months (May–Sep 2025).

COST

50x

Median inference-price decline per year, at fixed capability (Epoch AI).

TOOLING

MCP ships

Open protocol, 10,000+ servers, Linux Foundation (Dec 2025).

Sources: Anthropic, OpenAI, Google model cards (2025); Epoch AI, "LLM inference price trends," March 2025; Model Context Protocol release notes, 2024–2025.

THE TOOLING CURVE

MCP standardized connections between models and tools. Claude Code and Codex CLI made terminal-native orchestration robust enough for production use. Agent frameworks matured. The missing infrastructure for deploying agents has, in the space of eighteen months, stopped being missing.

Against those three tailwinds, the workforce signal is unmistakable. **Shopify CEO Tobi Lütke's** April 2025 memo made AI usage a baseline expectation and required teams to justify headcount requests against AI alternatives. **Amazon CEO Andy Jassy's** June 2025 memo told em-

ployees to expect a smaller corporate workforce as AI efficiency gains compound. **Salesforce's Marc Benioff**, on the firm's February 2025 earnings call, said Salesforce would hire no new engineers in 2025 against a 30% productivity gain.

These are not replacement stories in disguise. **Klarna's CEO** reversed course in May 2025 after concluding AI-only customer service produced lower-quality support, and resumed hiring human agents. **Duolingo's CEO** publicly reframed his company's "AI-first" memo as augmentation, not substitution, a month after issuing it. The firms that treated AI as a replacement broke things. The firms that treated it as amplification compounded their output.

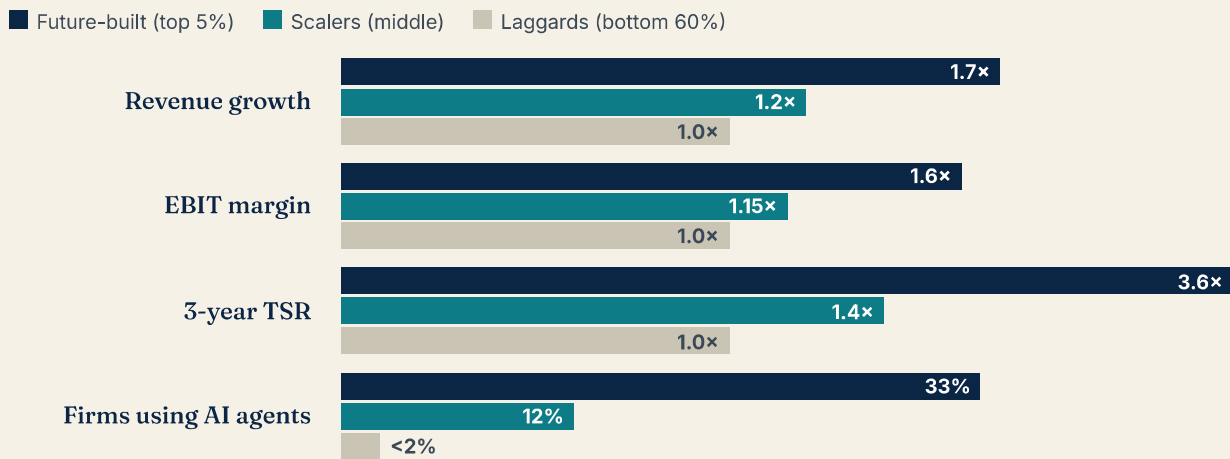
§ 04 — THE GAP The cost of inaction is no longer abstract

Conservative projections are now enough to be alarming. BCG's *Build for the Future 2025* study, based on 1,250 senior executives across nine industries, classified firms into three bands. The top 5% — "future-built" — are pulling away. They report **1.7× the revenue growth**, **1.6× the EBIT margin**, and **3.6× the three-year total shareholder return** of the 60% of firms classified as laggards.

They plan to spend more than twice as much on AI in the coming cycle.

Crucially for our thesis: **33% of future-built firms use AI agents. About 12% of middle-tier "scalars" do. Almost none of the laggards do.** The agent gap is not a lagging indicator of the value gap. It is a leading one.

EXHIBIT D The widening gap between AI leaders and the rest.



Source: BCG, *Build for the Future 2025*, September 30, 2025 (N=1,250 senior executives, 9 industries). Laggards indexed to 1.0×

McKinsey's *State of AI 2025* corroborates the pattern from the other side. About 6% of respondents qualify as AI high performers (attributing at least 5% of EBIT to AI). These firms are **2.8× more likely to have fundamentally redesigned workflows** (55% vs. 20%) and far more likely to have defined human-in-the-loop validation (65% vs. 23%). Redesign, not adoption, is where the money lives.

SIDEBAR · WHY THIS MATTERS ESPECIALLY FOR MID-MARKET

Margin erosion, not headline disruption

The specific risk for mid-market firms is subtler than loud disruption. It is margin erosion. Professional services firms grew revenue 4.6% year-over-year in 2025 — half the ten-year average — while EBITDA fell 36% (SPI Research, 2025). The Harvard Business Review has described the "flattening of the pyramid": AI is automating the junior-analyst work that funded the traditional leverage model. A competitor that deploys an agentic layer does not disrupt you loudly. It quietly prices 15–25% below your quote, delivers in half the time, and keeps the margin you surrendered.

//

The agent gap is not a lagging indicator of the value gap. It is a leading one.

§ 05 — CASE STUDY **MidCo — an 80-person B2B services firm**

MidCo is a composite, anonymized 80-employee specialty B2B services firm — roughly \$18M in annual revenue, average contract value \$55,000, 3 SDRs and 7 AEs, pipeline coverage 3.5× of quota, sales cycle 60–75 days. Its numbers are drawn from published benchmarks (Gong, Salesforce, Instantly, Gradient Works, Operatix) and reflect a defensible mid-market baseline.

Before

Sales and marketing ran on four tools that did not talk to each other: an outbound sending platform, a LinkedIn automation, a data enrichment service, and a CRM. Two of the three SDRs spent **roughly 60% of their time on research, data cleanup, and personalization** — consistent

with Salesforce's finding that reps spend only 28–30% of their week actually selling. Cold email reply rate averaged 3.8% (in line with Instantly's 2026 benchmark of 3.43%). Meetings booked per SDR per month sat at 15. Quarterly qualified pipeline was roughly \$1.6M. ChatGPT was used ad hoc by most of the team, unmonitored, in personal accounts.

The deployment — ten weeks

Weeks 1–2: mapping of current SaaS and AI usage, identification of a single high-leverage workflow. **Weeks 3–4:** pilot design and context assembly — CRM, calendar, inbox, and LinkedIn data unified through the agentic runtime. **Weeks 5–8:** three specialist agents deployed — email orchestration, LinkedIn prospecting, enrichment with anonymous website visitor identification. **Weeks 9–10:** human-in-the-loop review flows, dashboards, approval gates in place.


EXHIBIT E MidCo, before and after — same team, measurably more output (day 90).

■ Before ■ After (day 90)


Meetings / SDR / month



Cold email reply rate




Research time / account




SDR time spent selling



Qualified pipeline / quarter



Blended CAC



Source: Linkenite composite case study. Benchmarks: Gong State of Revenue 2025; Salesforce State of Sales 2026; Smartlead Cold Email Benchmarks 2026; Benchmarkit 2025; Operatix/Crunchbase SDR Metrics.

The delta was not won by higher send volume. Deliverability pressure on cold email has worsened every year since 2019. **The win came from quality of context.** The email agent knew who had replied to a similar thread six months earlier, which AE owned the account, which

meeting transcript revealed a buying signal, and which enrichment fields were missing. Humans still approved every send. They were amplified by a system that had mobilized the organization's own memory.

§ 05 CONT. — RESULTS TABLE & ECONOMICS

Same team, same headcount, measurably more output

TABLE 01 MidCo — 90-day results against mid-market benchmarks.

METRIC	BEFORE	AFTER (DAY 90)	DELTA	BENCHMARK SOURCE
Meetings / SDR / month	15	22	+47%	Gong State of Revenue 2025
Cold email reply rate	3.8%	6.5%	+71%	Smartlead 2026; Hunter.io
Research time per account	18 min	4 min	-78%	Salesforce State of Sales 2026
SDR time spent selling	28%	42%	+14 pp	Salesforce 2023–2026
Qualified pipeline / quarter	\$1.6M	\$2.1M	+31%	Agentforce benchmarks
Blended CAC	baseline	-18%	-18%	Benchmarkit 2025
Headcount	3 SDR · 7 AE	3 SDR · 7 AE	unchanged	—

Linkenite composite case study; benchmarks per source column.

Extrapolated conservatively: if MidCo holds the 31% pipeline gain and a 75% AE win rate, year-one incremental revenue is roughly **\$1.0–1.4M** against an AI tooling spend of well under \$100K. Payback is visible inside the first quarter.

SIDEBAR · WHY SALES AND MARKETING FIRST?

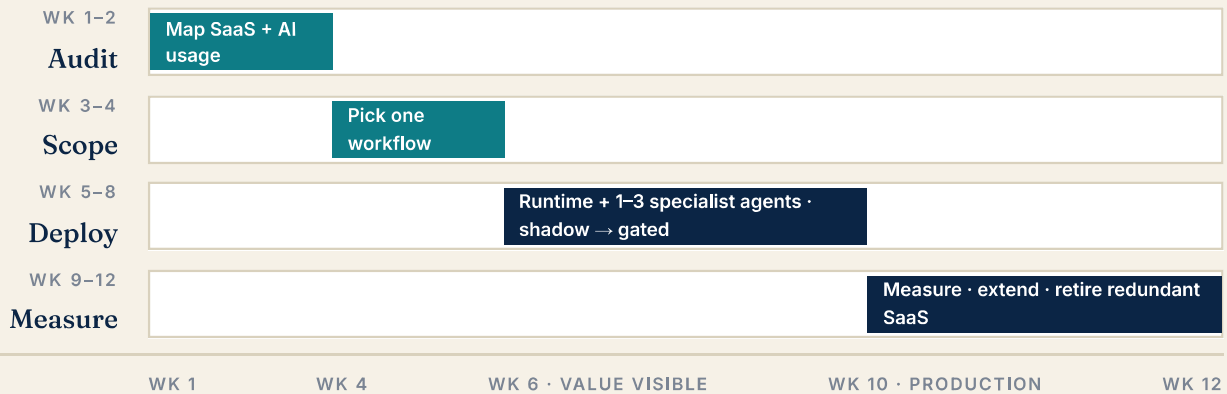
Outbound is the easiest place for most mid-market firms to prove the agentic layer. The feedback loops are short (days, not quarters), the data is structured (CRM, email, calendar), and the unit of work (a personalized, contextual touch) is both high-volume and high-leverage. Once the pattern is established, it migrates naturally into operations, support, and finance.

§ 06 — THE ROADMAP How to start — a 90-day plan

The mid-market's structural advantage is speed. MIT NANDA's 2025 data shows that mid-market firms that buy (rather than build) their AI capability reach production in

roughly 90 days. Enterprises typically need nine months or more. The reason is not talent; it is coordination overhead.

EXHIBIT F The 90-day roadmap — value visible by week six, production by week ten.



Linkenite deployment methodology; cross-referenced with MIT NANDA, *State of AI in Business*, July 2025.

Weeks 1-2 · Audit

Map every AI tool in use (official and unsanctioned), every SaaS the revenue team touches, every place organizational context lives. Identify the one or two workflows where cycle time and quality are bottlenecked by context gathering — almost always outbound, account research, or proposal drafting.

Weeks 3-4 · Scope the first workflow

Pick one end-to-end flow with a measurable business metric. Define the human-in-the-loop checkpoints upfront. Assemble the context sources. Agree on what "good" looks like.

Weeks 5-8 · Deploy

Stand up the orchestration runtime. Connect the data layer via MCP. Build one to three specialist agents against the chosen workflow. Run them alongside the human team in shadow mode first, then with approval gates, then with narrower gates as trust accrues. Measure against the baseline daily.

Weeks 9-12 · Measure, expand, retire

Publish the 90-day results. Extend to the adjacent workflow. Begin retiring redundant SaaS whose UI the agent no longer needs — a database view will do. Lock in a governance cadence. Value should be visible by week six. Production should hold by week ten. If either milestone slips, the workflow chosen was too broad — narrow and re-scope rather than extend the timeline.

§ 07 — PITFALLS Five failure modes that cancel most agentic projects

Gartner forecasts that over 40% of agentic AI projects will be canceled by the end of 2027, citing cost creep, unclear

value, and weak risk controls. The pattern is well understood and avoidable.

EXHIBIT G Five pitfalls — and the antidote to each.

PITFALL 01

Starting with everything

One workflow. One metric. Narrow or scope collapses.

PITFALL 02

Skipping the human loop

Approval gates are how the agent learns — and how trust accrues.

PITFALL 03

Replacing experts

+12–40% within frontier, –19pp outside it. Amplify, don't substitute.

PITFALL 04

Thin context

Context isn't a feature. It is the product.

PITFALL 05

Model lock-in

The frontier rotates quarterly. Stay model-agnostic.

Linkenite analysis; Gartner, June 2025; BCG/Harvard "Jagged Technological Frontier," 2023.

Starting with everything. Teams agenticize the whole business. Scope collapses. Focus on one workflow with one measurable outcome. **Skipping the human loop.** Klarna's 2025 walkback is the cautionary tale: AI replacing a customer-facing human with no editor above it degrades brand quality. **Replacing domain experts.** BCG's "jagged frontier" study found AI lifts consultants 12–40% *within* capability — and drops them 19pp outside it.

Underestimating context engineering. The difference between a useful email agent and a useless one is whether it knows the last three emails you exchanged with the prospect. **Over-indexing on one model provider.** Capability leadership rotates. The orchestration layer should be model-agnostic so you can route the right task to the right model — and switch without rebuilding.

§ 08 — OUTLOOK What this means for the next eighteen months

Three conclusions are reasonable from the current data. **First, every organization will have an agentic layer within two to three years.** The question is whether you build yours deliberately or inherit one through shadow adoption. Menlo Ventures tracked enterprise AI investment tripling from \$11.5B in 2024 to \$37B in 2025 (+222% YoY). That curve does not flatten in 2026.

Second, the mid-market window is open now and closes quickly. The capability curve is convex upward; the cost curve is convex downward. A six-month lead compounds into a multi-year one because the team that deploys ear-

lier collects the data, refines the context, and accumulates the playbooks that subsequent competitors will not easily replicate.

Third, the winners will treat this as an operating-model change, not a tooling purchase. McKinsey's high performers are 2.8x more likely to have redesigned workflows. BCG's future-built firms allocate 15% of AI spend to agents. Bain's 2025 *Technology Report* finds that AI delivers 10–15% productivity gains when bolted on, and 25–30% when workflows are redesigned. The math is not subtle.

“

The leaders of this transition will not be the firms with the most AI tools. They will be the firms whose agents know the most about the business.

An agentic engineering company.

[ABOUT LINKENITE](#)

The thesis

Linkenite designs, builds, and operates organizational agentic layers for mid-market companies — with an initial focus on the revenue workflow: outbound, prospecting, enrichment, and pipeline operations.

Your team already has the taste, the judgment, and the relationships. What it lacks is a layer that mobilizes the organization's own context on its behalf. **That is what Linkenite builds.**

We deploy in 90 days, work alongside your existing team, and instrument every step against the baseline you already measure. We stay model-agnostic at the orchestration layer so capability upgrades reach you without rebuilds.

[BOOK A 30-MINUTE CALL →](#)

The product stack — specialist agents, one runtime

REVENUE · EMAIL

InboxIn

Email orchestration and response agent. Drafts, routes, and follows up with full CRM and calendar context.

REVENUE · LINKEDIN

Approachin

LinkedIn prospecting agent. Context-aware connection, multi-touch sequences, reply triage.

REVENUE · DATA

SalesEnrich

Data enrichment plus anonymous website visitor identification — the missing input for every outbound decision.

■ Linkenite

Linkenite Oy · Helsinki, Finland

An agentic engineering company. We build organizational agentic layers for mid-market teams who would rather compound than sprint.

PERSPECTIVES

No. 01 — The Agentic Layer

No. 02 — The Context Audit (coming)

No. 03 — The 90-Day Playbook (coming)

CONNECT

info@linkenite.com

[LinkedIn](#)

[linkenite.com](#)

© 2026 Linkenite Oy. All rights reserved.
Published April 2026 · Helsinki.