



DOCUMENT

# The Operational Worker in the Three-Threads Transition

---

---

AUTHOR

LATENTTI

DATE

2026-05-02

*Draw the line before it forms by accident.*

# Contents

1. Why this matters now	3
2. The skills framework, adapted	5
3. Six people, two snapshots	7
4. How to lead this transition	17
5. The bimodal reality	19
5b. The augmentation differential inside one organisation	20
6. Watch-list — three signals individuals can read	22
7. Closing note	23
Glossary	24

# The Operational Worker in the Three-Threads Transition

## A companion to the Future of Enterprise Architecture 2026–2031

The main engagement framed the firm, the leader, the technical builder, and the architect. This companion zooms in on the executor — the sales person, the application developer, the customer service worker, the cashier, the warehouse picker, the back-office claims handler. It is their week, their identity, their cognitive turn that decides whether the bifurcation thesis from the board briefing shows up in workforce data. Six people, two snapshots each. The story behind the percentages.

---

## 1. Why this matters now

---

The board briefing argued that the dominant shape of 2026–2031 is bifurcation — a frontier cohort compounding integration depth, a long tail running the same software with shallower workflow change, a structural separation opening between them. That argument lives or dies in the workforce. The integration gap — 84% of developers using AI tools, only 44% with AI integrated into the actual workflow, 60% trust and falling — is not a procurement problem. It is a worker-level constraint. The Copilot is on the desk; the workflow has not been redesigned around it; the worker still describes themselves the same way they did in 2024. The thirty-point gap between adoption and integration is the gap between a tool and an identity.

The Operations Manager analysis put a number on what happens when that gap closes for one cohort and persists for another. Routine-heavy operational functions — customer operations, back-office finance, claims, support, scheduling — compress 45 to 60 percent in full-time-equivalent terms over five years in firms that complete the operating-model rewrite. That is a forecast about the next five Tuesday mornings of millions of workers. The bank teller, call-centre, and travel-agent reference classes show roughly thirty to fifty percent navigating to a new self-concept, twenty to thirty-five percent leaving the field, and a residual cohort in narrowed niches.

The HR Advisor's through-line was the deepest finding of the engagement. The cliff is on tool identity and output identity. The continuity is on discipline identity. A worker whose self-worth sits in the spreadsheet, the SOP console, the ticket queue, or the lines-of-code count faces compression of those things. A worker whose self-worth sits in the underlying craft —

judgment, customer relationship, trust, the discipline of getting the hard call right — finds that craft elevated. The pattern is consistent across mainframe, client-server, web, cloud, and SaaS.

This document gives that finding a face. Six workers, told as people rather than skill matrices. Some make the turn cleanly. Some struggle. One leaves the field. That distribution is not a writing choice; it is the bimodal reality the reference classes have already shown.

## 2. The skills framework, adapted

The Open Group's Architecture Skills Framework (G198, 2019) defines architect roles, skills, and depth of knowledge through seven categories crossed with four proficiency levels. It is a thinking tool, not a job-description template — which is why it generalises. A scaffold that worked for a self-conscious discipline in 2019 can carry the executor side of the same evolution in 2026.

We adapt G198 in two small ways: reduce seven categories to six recast for the executor, and keep the four proficiency levels verbatim.

#	CATEGORY	WHAT IT MEANS FOR THE WORKER	CONTINUITY
1	Domain Craft	The core thing this worker does — sales-craft, dev-craft, customer-service-craft, retail-craft, production-craft, claims-craft. The discipline-identity anchor.	The continuity layer. This is what survives.
2	AI Collaboration	Prompting, structured intent, validating outputs, recognising hallucination, knowing when to override, working with assistants effectively.	New. Level 1 in 2026, Level 3 expected by 2031 for most knowledge workers.
3	Agent Orchestration	Supervising one or more AI agents working on the worker's behalf — setting goals, defining guardrails, reviewing outputs, the agent-boss skill at the individual-contributor level.	New. Level 1 today, Level 2 to 3 expected by 2031 for many roles.

#	CATEGORY	WHAT IT MEANS FOR THE WORKER	CONTINUITY
4	Knowledge Stratification	Using personal-tier AI memory, contributing to team-tier knowledge, drawing from corporate-tier, respecting governance boundaries, knowing what is private, what is shared, what is public.	New. Emerging discipline.
5	Judgment and Exception Handling	What humans do in the bridge between deterministic rules and agent-mediated work — deciding when to escalate, handling the unscriptable.	Amplified. Level 2 today for many, Level 3 expected by 2031.
6	Trust and Relationship	Customer trust, peer trust, employer trust, AI accountability — the discipline-identity skills that survive the cliff.	Amplified. This is where the worker's value increasingly sits.

*Proficiency levels (from GI98, 2019, retained verbatim)*

*Level 1 – Background. Not a required skill, although awareness of the subject area is useful.*

*Level 2 – Awareness. Understands the background, issues, and implications sufficiently to be able to understand how to proceed further and advise accordingly.*

*Level 3 – Knowledge. Detailed knowledge of the subject area and capable of providing professional advice and integrating capability into architecture or design.*

*Level 4 – Expert. Extensive practical experience and applied knowledge in the subject area.*

The skill matrix at the end of each persona uses these six categories and four levels, applied to one person five years apart.

### 3. Six people, two snapshots

---

Persona 1 — Sami, application developer, 35, Finnish SaaS scale-up

#### **TODAY, 2026 — A TUESDAY IN AUTUMN**

Sami's day starts at half past eight with a black coffee and yesterday's pull-request queue. He opens VS Code, Copilot greets him in the sidebar, and he pairs with it for ninety minutes on a payments-edge feature. He accepts about a third of its suggestions cleanly, rewrites another third, ignores the rest. The hard part of the morning is not typing — it is getting the specification clean enough that the agent stops drafting in the wrong direction. After lunch he reviews two pull requests from a junior who joined six months ago and has never written software without an AI assistant; the code looks plausible but does not handle a webhook retry cleanly, and he writes the comment carefully. He leaves at six having shipped less than he wanted and feeling, as he often does now, that the typing got faster while the thinking got harder. He still believes the job is to ship correct software. He does not yet describe himself as anything other than a developer.

#### **BETWEEN NOW AND 2031 — WHAT CHANGES**

Late 2027 is the inflection. The team adopts an internal agent platform that builds whole bounded features against a written specification. Sami's first reaction is that the specification is the bottleneck — and writing it well is harder than writing the code was. By mid-2028 his calendar has flipped: less time in the editor, more in design reviews, more reading agent-generated code than authoring it, and a new kind of work where he tunes the agent's behaviour the way he used to tune database indexes. He picks up the verifier-engineer label from a Stockholm conference and starts using it. The cognitive turn is uncomfortable for about eight months in 2029. He misses the typing.

#### **2031 — A TUESDAY IN AUTUMN, FIVE YEARS ON**

Sami arrives at nine and the pod's three agents have already produced overnight a candidate implementation of yesterday's specification, a verification harness, and three escalation flags. He spends the morning on the third flag — the agent has been confused by a regulatory edge case in a Norwegian statute, and Sami calls a contact at the national supervisor to confirm interpretation before shipping. At two he runs a design session with the new juniors, who do not write foundation code in the editor but learn to specify and verify in a sandboxed environment that deliberately restricts agent access for the first eighteen months. He still calls himself an engineer. He now describes engineering as judgment about systems rather than the production of code.

## WHAT SHIFTED IN HIS SKILL PROFILE

SKILL CATEGORY	2026	2031	NOTE
Domain Craft	4	4	Engineering rigour deepens, not narrows
AI Collaboration	2	4	Daily craft, not adjunct
Agent Orchestration	1	3	Sets guardrails for a pod of agents
Knowledge Stratification	1	3	Curates team-tier engineering memory
Judgment and Exception Handling	3	4	The flagged-and-escalated cases become the day
Trust and Relationship	2	3	Trust in the verification chain becomes load-bearing

### Persona 2 – Anna, B2B account executive, 32, Helsinki

#### TODAY, 2026 — A TUESDAY IN AUTUMN

Anna takes her first call at half past nine with a procurement lead who has been ghosting her for two weeks. She has prepared with a Glean-style search that surfaced internal notes, the customer's recent press releases, and a summary of eleven email threads. She catches one error before the call. The call goes well; she books a follow-up demo for Friday. Between calls she edits Copilot-drafted follow-up emails, accepting half of the language and rewriting the openings because the AI's voice is bland and a little too eager. The afternoon is two demos and the weekly forecast call. Anna's quota is up six percent year over year and she is hitting it. The job, as she describes it to her parents, is still selling — finding the customer, understanding their problem, building trust, closing the deal.

#### BETWEEN NOW AND 2031 — WHAT CHANGES

In 2027 the firm rolls out an outbound agent that researches accounts and drafts personalised first-touch emails at volumes Anna could never have hand-written. Her job shifts from discovery to qualification and close. She resists for a quarter, then the autumn 2027 results show her close rate on agent-sourced opportunities is two-and-a-half times the team average

— her ear for which procurement leads are real is the part the agent cannot do. By 2029 she is what the firm calls a deal-architect, running multi-stakeholder pursuits where the agent layer handles briefing, scheduling, and follow-up, and her time is in the room with the customer.

**2031 — A TUESDAY IN AUTUMN, FIVE YEARS ON**

Anna's morning is one ninety-minute conversation with a CFO at a Norwegian customer, prepared by an agent that has summarised every internal touchpoint, flagged three procurement-policy clauses that will matter, and surfaced a competitor product-release the CFO has probably read. The conversation is harder than any she had in 2026 because the customer arrived more prepared, briefed by agents of their own. She handles a regulatory question by calling the customer's compliance lead directly — a relationship she has built over four years. Her quota is denominated differently now: not deals closed, but accounts retained and expanded, with a trust-incident threshold that claws back commission if breached. She still calls herself a sales person. She also says, when pressed, that she is the trust layer between her firm and her four largest customers.

**WHAT SHIFTED IN HER SKILL PROFILE**

SKILL CATEGORY	2026	2031	NOTE
Domain Craft	3	4	Selling deepens into trust-architecture
AI Collaboration	2	3	Briefing and drafting handled fluently
Agent Orchestration	1	3	Runs an outbound and follow-up agent pod
Knowledge Stratification	2	3	Account memory contributed to team-tier
Judgment and Exception Handling	3	4	Reads which procurement signals are real
Trust and Relationship	4	4	Already her core; now the explicit metric

## Persona 3 — Jukka, customer service representative, 28, Nordic bank contact centre

### TODAY, 2026 — A TUESDAY IN AUTUMN

Jukka logs in at eight, headset on, and works fifty to eighty interactions across an eight-hour shift. The platform suggests responses; he accepts most for routine cases and overrides about a fifth. The cases he likes are the difficult ones — the elderly customer who has been scammed and is panicking, the small-business owner whose payroll bounced, the customer disputing a six-month-old fraud charge who has not been listened to by three previous reps. He is good at those. His team lead has started calling him in for the toughest escalations. Jukka still describes himself as a customer service rep, the way his older brother does after fifteen years on the job. He has been on the team for three years and is starting to wonder how much longer this role exists.

### BETWEEN NOW AND 2031 — WHAT CHANGES

By late 2027 the bank has shifted tier-zero and most of tier-one to autonomous agent handling. Jukka moves to a smaller pool of exception-handling specialists. About forty percent of the original team is reassigned, retrained for in-branch advisory, or leaves. Jukka stays. His training shifts to judgment-and-trust scenarios — recognising vulnerable customers, handling regulator-triggering cases, calming customers whose first three interactions with the agent layer made things worse. In 2028 the bank formalises his role as a customer-trust specialist with a salary band noticeably above the old representative band. In 2029 a new manager pushes metrics that reward speed rather than quality on exceptions. Jukka pushes back. The 2030 retrospective shows trust incidents on his team below the firm average; the framework is revised.

### 2031 — A TUESDAY IN AUTUMN, FIVE YEARS ON

Jukka takes twelve interactions in his shift. Each has been escalated by an agent that flagged it as outside its confidence band — a vulnerable customer, a complex multi-product complaint, a regulatory edge case, a customer whose anger needs to be heard before anything else can happen. He spends most of his time listening. He documents each case carefully; the documentation becomes part of the team-tier memory the agent layer will draw on next week. At four he runs a coaching session with two newer specialists. He earns thirty percent more than he did in 2026. He still calls himself a customer service person. He says, when his daughter asks, that his job is to be the human voice when the system has not been good enough to a customer.

## WHAT SHIFTED IN HIS SKILL PROFILE

SKILL CATEGORY	2026	2031	NOTE
Domain Craft	3	4	Service deepens into trust craft
AI Collaboration	2	3	Drafts, summaries, pre-call briefing
Agent Orchestration	1	2	Reads what the agent flagged and why
Knowledge Stratification	1	3	Documents exception reasoning into team tier
Judgment and Exception Handling	3	4	The whole job is now this
Trust and Relationship	3	4	The discipline anchor

### Persona 4 – Maria, supermarket cashier, 45, Finnish retail chain

#### TODAY, 2026 — A TUESDAY IN AUTUMN

Maria starts the morning shift at six. The store has eight lanes; four self-service, four staffed, and the self-service share is up from sixty to seventy percent over the last two years. Maria works lane three. She knows twenty regular customers by name. She is good at the small things — finding a product code for an item with a missing barcode, dealing with a frozen payment card, holding the queue together when a self-service lane crashes. Her shift is interrupted three times today to help at the assisted-checkout area. The store manager has mentioned that the next refit will increase self-service further. Maria has been with the chain for twenty years and is the union representative on her shift. She is starting to feel, without quite saying so, that the store has not yet decided what the cashier role is for.

#### BETWEEN NOW AND 2031 — WHAT CHANGES

The 2027 refit removes two staffed lanes for agent-mediated assisted-checkout zones. Maria is reassigned to assisted-checkout supervisor — marketed internally as a promotion. She does not take it as one for six months. In late 2028 the chain pilots an in-store advisory role: a senior staff member roams the aisles with a tablet, helps customers find products, runs demonstrations, manages click-and-collect. Maria volunteers and is selected. The first eighteen months are hard. She learns the tablet and the inventory system; her pay band is

reset. She considers leaving twice. She stays because her colleague Helena leaves and does not find another retail job — ends up on a short-term contract at a logistics warehouse. By 2030 Maria is one of three in-store advisors, runs click-and-collect, and trains new staff. She is paid more than she was as a cashier. The eighteen-month gap was real.

### 2031 — A TUESDAY IN AUTUMN, FIVE YEARS ON

Maria opens the store at seven, walks the aisles with a tablet, checks three click-and-collect orders before customers arrive. At ten she runs a coffee-machine demonstration and has eleven customers around her by the end. She helps an elderly couple find a product the agent in their phone has wrongly told them is in stock; she finds it, helps them check out, notes the inventory error so the system can be corrected. The supermarket has roughly the same total revenue as in 2026 with about sixty percent of the staff. Maria is one of the ones who stayed. She does not call herself a cashier any more. She is a store specialist, the chain's term, and the Finnish word sounds much like what she would have used to describe her aunt who ran a small grocery in the seventies — a person who knew the customers, the products, and the rhythm of the season.

### WHAT SHIFTED IN HER SKILL PROFILE

SKILL CATEGORY	2026	2031	NOTE
Domain Craft	3	4	Cashier-craft replaced by store-craft
AI Collaboration	1	2	Tablet-mediated inventory and tasks
Agent Orchestration	1	2	Supervises self-checkout zone
Knowledge Stratification	1	2	Reports inventory exceptions back to system
Judgment and Exception Handling	2	3	Customer-facing edge cases
Trust and Relationship	3	4	Now the explicit value of the role

## Persona 5 – Petri, warehouse fulfilment worker, 38, e-commerce centre

### TODAY, 2026 — A TUESDAY IN AUTUMN

Petri is a picker-packer on the morning shift. He wears a hand scanner, follows the route the warehouse-management system gives him, picks, packs, drops items on the conveyor. He hits his pick-rate target most days. The warehouse runs an automated robot fleet on high-velocity SKUs; Petri's section — slower-moving, oddly-shaped items — is still mostly human. He earns a reasonable hourly wage with overtime, has worked here for nine years, has a daughter at school nearby. The supervisor talks about a major system upgrade next year. Petri does not pay close attention. The job is what it is; the hand scanner tells him where to go.

### BETWEEN NOW AND 2031 — WHAT CHANGES

The 2028 refit is the inflection. The operator brings in a second-generation robotic fleet that handles ninety percent of the SKUs Petri's section used to cover. The human role at his station is reduced to robot supervision, exception handling on damaged or mis-routed items, quality checks on outbound shipments. Site headcount drops forty percent over eighteen months. Petri keeps his job through the first round but is reassigned to a multi-skilled role across three stations on rotating shifts. His take-home pay drops fifteen percent because overtime evaporates. He retrains on the robot-supervision interface and finds it frustrating. In 2030 a second consolidation is announced. He is offered a transfer to a regional hub forty kilometres away or a severance package. The hub is not feasible because of his daughter's school. He takes the severance. The co-funded retraining package covers a six-month course in heating-and-cooling installation, a skills-shortage trade in his region. He starts an apprenticeship in early 2031.

### 2031 — A TUESDAY IN AUTUMN, FIVE YEARS ON

Petri is twelve weeks into his apprenticeship as a heat-pump installer. The trade is busy; regional subsidies are pushing residential heating-system replacements, and his contractor has more work than installers. He earns less than at the warehouse and works longer hours. When he qualifies, he thinks the work will pay better than the warehouse did, and the trade is hard to automate because it happens in people's basements with bespoke pipework. The eighteen months at the warehouse before the severance, watching the robots take over his work, were the worst of his career. He does not regret leaving. He is wary about whether he has bought himself a decade or longer before the next compression. He is honest about not knowing.

## WHAT SHIFTED IN HIS SKILL PROFILE

SKILL CATEGORY	2026	2031	NOTE
Domain Craft	3	2	Pivoted; new craft, low level, growing
AI Collaboration	1	1	Trade work less AI-mediated for now
Agent Orchestration	1	1	Not relevant in current role
Knowledge Stratification	1	1	Not relevant in current role
Judgment and Exception Handling	2	3	Trade decisions on bespoke installations
Trust and Relationship	2	3	Customer trust in their home matters

### Persona 6 – Liisa, insurance claims handler, 50, Nordic insurer

#### TODAY, 2026 — A TUESDAY IN AUTUMN

Liisa has been handling motor and household claims for twenty-five years at the same Nordic insurer. She is the senior on her team, the one juniors come to when a claim is awkward, the one regulator information requests are routed through. She handles fifteen to twenty open claims at any moment. The firm rolled out a claims-triage assistant last year that scores incoming claims for fraud risk and suggests settlement bands. Liisa overrides it about a third of the time. Her team lead notes the override rate in her monthly review. Liisa thinks the assistant is wrong on the cases that matter most — the ambiguous ones — and right on the cases that are easy enough that any junior could handle them anyway. She is two years from an early-retirement option, undecided. She still describes the job exactly the way she did in 2002.

#### BETWEEN NOW AND 2031 — WHAT CHANGES

The 2027 system upgrade is bigger than the 2026 assistant. The insurer moves to an agentic-claims platform: straightforward motor claims settled end-to-end by the agent layer with continuous-compliance audit, household claims under a settlement threshold agent-handled with human spot-check, and only genuinely complex claims arriving at a human handler. Liisa's caseload halves in volume but each case is two to three times harder. The cases on her

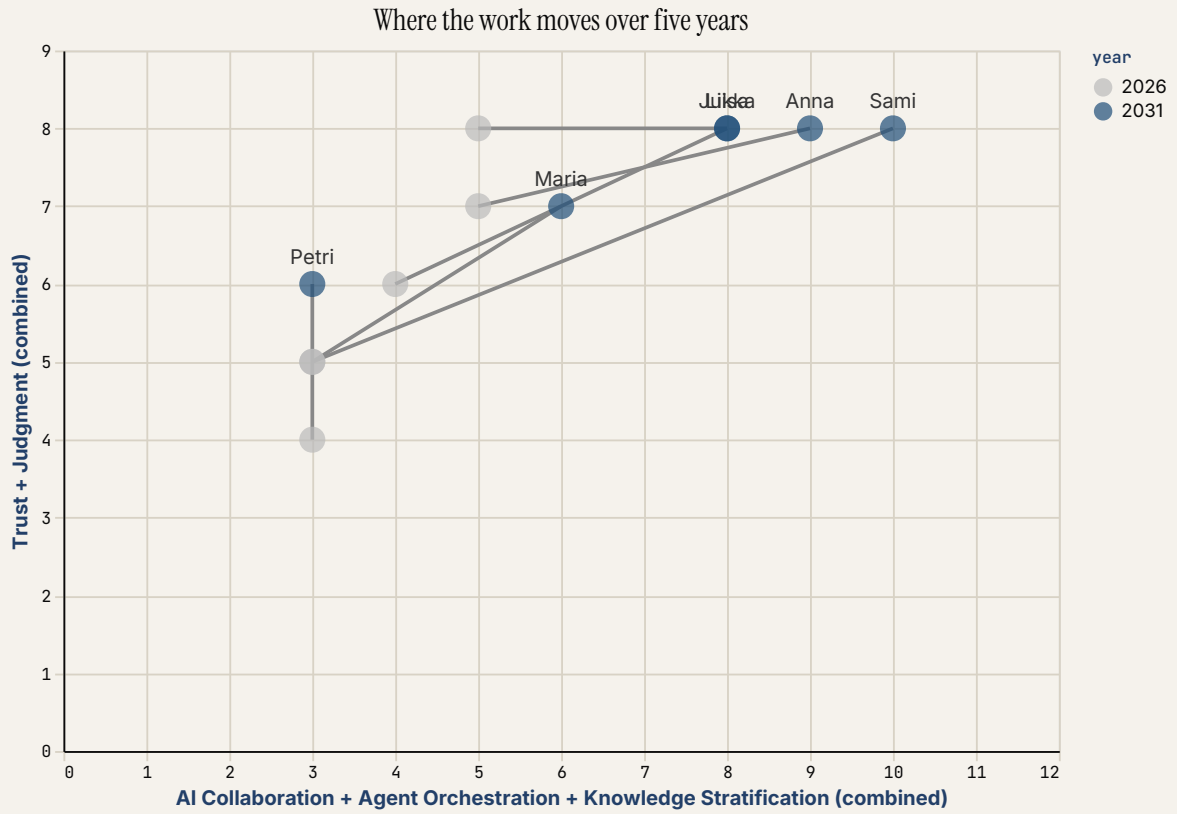
desk are uniformly the disputed, multi-party, regulator-triggering, emotionally difficult ones. She finds the cognitive load heavier than the old mixed caseload. In 2028 she requests a part-time arrangement; the firm grants it with a reset of her seniority track. In 2029 she leans into a role the firm has named claims-trust specialist — focused entirely on customer-facing trust and complaint, with the agent layer handling underwriting and settlement logic. She finds this better suited to her strengths. By 2031 she is content but ambivalent. The role she trained for in her twenties no longer exists; the role she does now is a different job she has had to make peace with at fifty-five.

### 2031 — A TUESDAY IN AUTUMN, FIVE YEARS ON

Liisa works three days a week as a claims-trust specialist. Today she handles four conversations: a grieving household claimant where the agent's tone was wrong; a motor claim with a regulator complaint attached; a four-party commercial dispute; a vulnerable-customer escalation flagged by the agent. She produces written reasoning on every decision because it becomes audit material and team-tier memory. On Wednesdays she trains a new specialist. She is paid less in absolute terms than she would have been on the old senior-handler track, more in real terms after inflation, and works fewer hours. She tells younger colleagues the cognitive turn took her three years and was harder than she expected.

### WHAT SHIFTED IN HER SKILL PROFILE

SKILL CATEGORY	2026	2031	NOTE
Domain Craft	4	4	Claims-craft retained, narrowed
AI Collaboration	2	3	Reads agent reasoning fluently now
Agent Orchestration	1	2	Reviews flagged cases, sets escalation patterns
Knowledge Stratification	2	3	Documents reasoning into team tier
Judgment and Exception Handling	4	4	Already her core; now the whole job
Trust and Relationship	3	4	The explicit value of the role



Five of six personas move toward higher trust-and-judgment scoring; Petri leaves the field and enters a new domain at low proficiency. The cliff is real for at least one of every six.

## 4. How to lead this transition

---

The transformation does not happen to workers; it happens to a firm with workers in it. Three phases over five years, each with a different organisational role.

### Phase 1 – Acclimate, 2026 to 2027

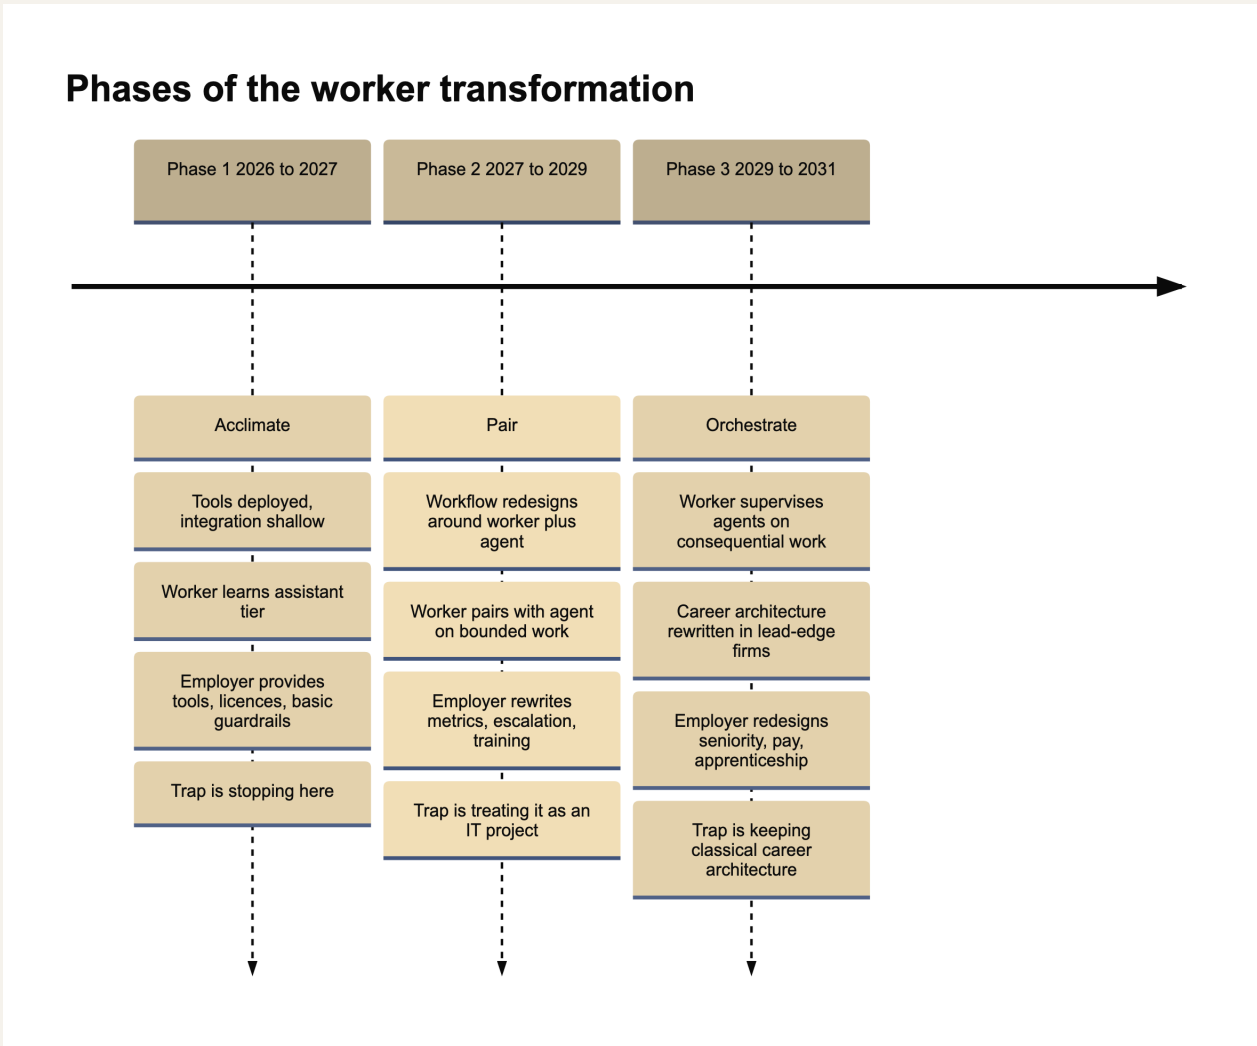
AI tools are widely available. Most workers have them on the desk. Integration is shallow. The worker's job is to learn the assistant tier — fluency, not mastery. The employer's job is to provide the tools, license them properly, train basic prompting, and set the first data guardrails. The trap is stopping here. Most firms will. The 84% / 44% / 60% numbers describe a workforce held in Phase 1 by an absence of operating-model investment. Workers in Phase 1 firms experience AI as a faster typewriter and do not yet describe their job differently than they did in 2024.

### Phase 2 – Pair, 2027 to 2029

Workers begin pairing with AI as a collaborator on bounded work, not just an assistant on individual tasks. Sami's specification-and-verification rhythm. Anna's qualification-and-close shift. Jukka's exception-only caseload. The workflow restructures around what the worker does and what the agent does. The employer's job is to redesign the workflow rather than bolt AI on top of the old one — rewriting metrics, escalation patterns, training, performance management. The trap is treating workflow redesign as an IT project rather than an operating-model rewrite. Firms that go in with a Copilot rollout plan and come out with their old org chart, their old career ladder, and their old performance-review template have not done Phase 2.

### Phase 3 – Orchestrate, 2029 to 2031

Workers supervise agents on consequential work. The agent-boss skill at the individual-contributor level. Anna's deal-architect role. Jukka's customer-trust specialist role. Liisa's claims-trust specialist role. Their value sits in judgment, exception handling, trust, and the discipline of the craft itself. The employer's job is the deepest of the three: career-architecture redesign. What does seniority mean when AI flattens output? What is fairly compensated? How does a junior become senior when an agent writes the foundation work? The trap is keeping classical career architectures and getting outcompeted by firms that have updated them. The career architecture is the operating-model artifact that takes longest to rewrite, and firms that do not rewrite it will lose senior-track talent to firms that do.



Three phases. Each requires a different organisational investment, and each has a specific failure mode that distinguishes lead-edge firms from the long tail.

## 5. The bimodal reality

---

The honest reading of the six personas is that they are not all the same story.

**Some workers make the turn cleanly.** Sami, Anna, and Jukka anchor on their domain craft — engineering rigour, customer trust, complex case judgment — and the AI layer becomes an amplifier. They had craft pride going in. The cognitive turn is real but bounded; the work in 2031 is more interesting and better-compensated than it was in 2026.

**Some workers struggle.** Maria's eighteen-month gap before the in-store advisor role stabilised was a real period of doubt and pay reset. Liisa's three-year cognitive turn at fifty was harder than she expected and required a part-time arrangement and a seniority reset. The bridge is built by employer, colleague, and worker together, and firms that under-invest in it lose the workers who would otherwise have made the turn. The bridge is not a training course. It is a sustained organisational commitment to the language, the rhythm, and the safety net of the transition.

**Some workers leave their field.** Petri's role compresses faster than the warehouse can absorb. He retrains into a trade with a current shortage. He may be fine; he may face the same compression in ten years. Mid-career operational specialists in routine-heavy work are the most exposed cohort. Reference data from prior platform shifts puts the leaving-the-field rate at roughly twenty to thirty-five percent of the affected workforce — millions of individuals across the period, regionally concentrated where routine-heavy operational employment was the dominant local labour market. Society, not just the employer, has a role here, in re-skilling infrastructure and transition support. Pretending this part is small is dishonest, and dishonesty about it makes the bifurcation thesis from the board briefing harder to act on, not easier.

## 5b. The augmentation differential inside one organisation

---

The bimodal pattern above is the exit-side story. There is a second, harder story playing out inside firms that retain their workforce. By 2031, two workers in the same role at the same firm can differ by an order of magnitude — sometimes two — in AI-augmented output. This is unprecedented intra-organisational productivity dispersion, and most firms are not designed for it.

Three archetypes coexist on the same floor.

**The 0x worker** is at pre-AI baseline. Output looks much as it did in 2024. Several reasons, often combined: skepticism about AI, discomfort with the tooling, role design that does not lend itself to augmentation, an employer that licensed the assistant tier but never restructured the workflow around it, or simply that nobody coached the worker through Phase 1 of the transition. Real and common — likely a quarter of the typical knowledge-work cohort even in firms that consider themselves AI-mature.

**The 10x worker** has folded AI into the daily workflow. Drafts come from an assistant; routine analysis goes to an agent; meetings produce searchable notes by themselves; the worker spends time on the parts that matter and the parts that need them. The realistic median for engaged knowledge workers by 2031 in firms that did the workflow redesign in Phase 2.

**The 100x worker** has restructured their work around agent supervision, personal AI memory, and multi-tier knowledge contribution. Several agents work on their behalf simultaneously; they invent the prompts, the validation patterns, and the team-tier conventions that the rest of the floor adopts six months later. Rare. Possibly five to ten percent of any large knowledge-work cohort, concentrated in roles where the work is highly augmentable. Sami is heading toward this band; Anna may reach it for proposal-and-pursuit work but not for relationship work; Maria's role has no 100x ceiling because the variance is bounded by the work itself, not the worker.

The implications a firm has to manage are not subtle.

**Compensation.** Pay-for-output is socially toxic at hundred-fold dispersion. A firm cannot pay one knowledge worker the salary of a hundred peers, and would not survive the cultural consequences if it tried. Pay attaches to role, judgment, accountability, and contribution-to-team-tier-knowledge — not to raw throughput. Firms that update compensation theory get the 100x workers and keep the 10x cohort engaged; firms that do not will see their best people leave for places where the gap is recognised differently — often as equity, autonomy, or scope, rather than salary multiples.

**Performance management.** The dispersion is real and the manager has to see it without using it as the appraisal axis. The 100× worker's value is partly in their effect on others — the patterns they establish, the conventions the team adopts, the corrections they push into team-tier knowledge. A manager who appraises only on individual output misses this and demotivates the worker doing the most durable work.

**Team cohesion.** The 0× worker watches the 100× worker and either resents or learns. The 100× worker watches the 0× worker and either supports or disengages. Healthy teams find rituals that bridge — the 100× worker explicitly teaches; the 0× worker is given the structured support to move toward 10×. Unhealthy teams ossify into castes. Most firms underestimate how fast this ossifies and how hard it is to undo.

**Hiring and internal mobility.** Identifying 100×-capable workers in advance is hard. There is no reliable test. The signals are indirect — comfort with ambiguity, willingness to teach the system, restlessness with their own previous baseline. Firms that build internal mobility paths so that 100× workers can change scope without leaving keep them; firms that do not lose them.

**Career architecture.** Classical career levels (junior / senior / staff / principal) compress because the output gap between adjacent levels collapses. Some firms respond by widening levels (one senior is now the work of three); others by introducing new tracks (an "orchestrator" track parallel to the IC track). Both patterns work. Holding the 2024 architecture and pretending the dispersion is not there does not work.

**The asymmetric-output ethic.** When one person's hour is plausibly worth ten or a hundred peers' hours, a firm has a partly philosophical question to answer: what is fairness now? Most firms will resolve this implicitly through compensation and role design rather than explicitly. The firms that resolve it explicitly — through stated compensation principles, through transparent career architecture, through honest communication about why the gap exists and how the firm chooses to handle it — keep the trust of both ends of the distribution. The firms that handle it through silence and ad-hoc bonuses lose the trust of both.

The augmentation differential is not a problem to be eliminated. It is a structural feature of the next five years and a permanent one beyond. The firms that name it, design for it, and lead through it will be the firms whose workforce composition matches the bifurcation thesis from the board briefing. The firms that pretend everyone is roughly the same will keep finding that their best people leave and their median people stall.

## 6. Watch-list — three signals individuals can read

---

A worker who wants to know which side of the cliff they are on does not need a certification. Three questions, asked honestly, about how the week actually feels.

**Are you spending more time on judgment, exceptions, and trust than a year ago, or less?** If the share of routine work in the week is rising relative to the harder, interpretive, interpersonally-weighted work, the worker is sliding toward the cliff. If the share of judgment work is rising, the worker is on the continuity side. The signal is at the level of the week, not the day; a six-month trend is the indicator.

**Are you teaching the system as much as it is teaching you?** Workers on the continuity side contribute corrections, conventions, and hard-case reasoning into the team-tier knowledge the agent layer draws on. Workers on the cliff side only consume — they accept or rewrite a draft and the rewrite goes nowhere. If the worker is not building durable team-tier value, they are running an individual productivity gain the firm cannot compound. Fine for the worker today; not fine for the worker over five years.

**Are you the one your colleagues come to when the AI gets it wrong?** The human-judgment-on-edge-cases role is amplified, not compressed. The worker colleagues escalate to is the worker whose discipline identity is load-bearing. The worker no one escalates to is the worker whose value the agent layer can substitute for at the next platform refresh. This is the cleanest signal and takes the most honesty to read.

These are not credentials. They are everyday signals. Honest answers tell a worker their trajectory; uncomfortable answers tell them what to work on.

## 7. Closing note

---

The change is happening at the speed of organisational redesign, not the speed of model capability. The Copilots and the agent platforms will be everywhere by 2027. The firms that translate ubiquity into productivity are the ones that rewrite workflows, metrics, training, performance management, and career architectures. The firms that do not will have the same tools and the same gap they have today. Workers who re-anchor on craft, judgment, and trust have the durable identity. Workers who hold to tool identity or output identity face the cliff. Employers who build the bridge get the workers who would otherwise have left. The bifurcation we see at the firm level is built person by person, week by week, conversation by conversation. The firm-level shape is, in the end, the sum of these weeks.

# Glossary

---

**Three threads.** The board briefing's framing of 2026–2031: agentic orchestration, knowledge stratification, and determinism routing.

**Integration gap.** The thirty-point gap between adoption (84% of developers using AI tools) and workflow integration (44%). Identity-level, not tooling-level.

**Bifurcation / two-speed enterprise.** The likely shape of 2031: a frontier cohort compounding integration depth, a long tail in shallow adoption, structural separation opening between them.

**Agent boss.** Term from Microsoft's 2025 Work Trend Index for someone who builds, delegates to, manages, and evaluates AI agents.

**Agent orchestration.** Supervising one or more AI agents on the worker's behalf — goals, guardrails, output review, escalation.

**Knowledge stratification.** Multi-tier knowledge architecture spanning personal, team, function, corporate, cross-organisation, and public tiers, each with its own permissions, governance, and decay rules.

**Determinism routing.** The deliberate choice of which surfaces stay rule-based and which move to AI agents. Rules win where audit trails, latency, and explainability matter; agents win where variance is high.

**Foundation model.** A large, generally-trained AI model (Anthropic, OpenAI, Google, Meta) specialised or prompted for specific tasks.

**MCP (Model Context Protocol).** Open protocol for AI model interaction with tools and context. Released by Anthropic in November 2024; donated to the Linux Foundation in December 2025.

**AI collaboration.** Working with AI assistants effectively — prompting, structured intent, validating outputs, recognising hallucination, knowing when to override.

**Judgment layer.** The bridge between deterministic rules and agent-mediated work where humans decide when to escalate and how to handle the unscriptable.

**Discipline identity vs. tool identity.** The HR Advisor's load-bearing distinction. Tool identity ("I run the SOP console," "I write the SQL") faces compression. Discipline identity ("I am the judgment layer for the customer outcome," "I steward the data layer") survives and is often elevated.

---

End of 17-worker-transformation.md.