Microsoft

Agentic Transformation Patterns Playbook

A practical guide to choosing, scaling, and operating AI agents across your organization

How to use this playbook

This playbook is designed for self-guided use — work through it at your own pace

1. Understand the landscape

Start with the Assist → Execute shift (next page) to understand why different agent initiatives need different operating models.

2. Identify your patterns

Section 1 covers six adoption patterns. Read the overviews, then deep-dive into the 1–2 patterns that match your organization's current priorities.

3. Assess your maturity

Section 2 introduces the maturity model. Use the detailed 5×5 diagnostic to assess where you stand today across five capability drivers.

4. Find your gap

Each pattern has a target maturity profile. Compare your current state to the target — the biggest gap is your scale-breaker.

5. Build your operating model

Section 3 introduces the Center of Excellence — the vehicle that closes your gaps and scales agents safely.

6. Start your 90-day play

Section 4 gives you a concrete plan. Pick your pattern, name your owner, find your scale-breaker. Start Monday.

The fundamental shift: Assist → Execute

Understanding this shift is the foundation for everything in this playbook

AI agents are moving from assisting humans to executing work. This single shift changes everything about how you govern, own, and operate agents — and it's why there is no "one size fits all" framework for all agent initiatives.

ASSIST

- → Agent supports human decision-making
- → Human decides what to do
- → Human executes the action
- → Human is fully accountable

Low risk • Familiar governance

⟹

EXECUTE

- → Agent performs work across systems
- → Agent acts on decisions
- → Agent orchestrates workflows
- → Human oversees outcomes

Higher risk • New operating model required

This shift creates four new demands:

Ownership

Who is accountable for this agent?

Risk

What happens when it goes wrong?

Lifecycle

Who maintains and improves it over time?

Governance

What is it allowed to do and not do?

Why adoption patterns matter

Not all agent initiatives are the same — treating them the same is what breaks scale

Organizations deploy agents for fundamentally different purposes. An agent that drafts emails for individuals is a completely different bet than an agent that processes insurance claims autonomously. They require different governance, different ownership, different success metrics, and different levels of organizational maturity.

Adoption patterns give you a classification system — a way to name what you're doing so you can match the right operating model to each initiative. Without this classification, organizations apply the same framework for everything, over-governing simple agents and under-governing complex ones.

Patterns are design choices, not stages:

→ You don't progress through them in order — you choose based on your intent

→ Most organizations pursue 2–3 patterns simultaneously

→ Each pattern demands a different depth of maturity across five capability drivers

→ Your pattern determines WHERE you need to invest — not just how much

→ Starting with the wrong pattern for your maturity level is a primary reason agents stall

The following pages introduce each pattern in detail — what it is, what it demands, and how to know if you're ready.

Section 1
Frontier Transformation Patterns

A transformation pattern is a design choice for how work runs with agents — who does the work, who decides, and how it's governed.

Frontier Transformation Patterns

Different agents, different rules. Frontier firms make this explicit by categorizing agent work into a small number of repeatable transformation patterns.

Enrich
employee experiences

Employee AI Enablement

Make people more capable

Every employee uses capable AI assistants

Different adoption model

Business Expert Empowerment

Scale SME judgment

Expert's knowledge scaled to the org.

Different ownership model

Reinvent
customer engagement

External Engagement

Customer and partner agents

Agents serving external stakeholders

Different trust boundary

Reshape
business processes

Workplace & IT Services

Run internal services

Self-service IT and workplace operations

Different service model

Core Business Processes

Transform workflows

Agents woven into business-critical flows

Different risk profile

Bend the curve
on innovation

AI-first Business Capabilities

Net new capabilities

Things that weren't possible before AI

Different everything

These are design choices, not stages. Most organizations run 2–3 simultaneously.

Employee AI Enablement

PATTERN 1

Employees use AI agents to research, analyze, draft, and automate personal workflows — while people remain accountable for all decisions and outcomes. This is the most accessible pattern and typically the starting point for most organizations.

WHAT AGENTS DO

Agents handles routine, time-consuming tasks: drafting content, summarizing documents, researching topics, scheduling, and automating personal workflows. They act as a capable assistant that amplify individual productivity.

WHAT THE HUMAN DOES

The human retains full decision-making authority. They review agent output, apply judgment, and are accountable for every action taken. The agent recommends; the human decides.

EXAMPLE USE CASES

- → Content drafting, editing, and summarization
- → Meeting preparation and follow-up automation
- → Personal research and information synthesis
- → Role-specific knowledge assistants
- → Email triage, scheduling, and workflow automation
- → Data analysis and report generation

OPERATING SHIFTS

People

Doing repetitive tasks → Making better decisions faster

Agents

In-app assistance → Embedded in daily workflows

Governance

Tool-use policies → Identity and data-bounded usage

Metrics

License usage → Output quality + decision speed + time saved

Employee AI Enablement: Maturity Profile

Target maturity across five capability drivers

TARGET MATURITY PROFILE

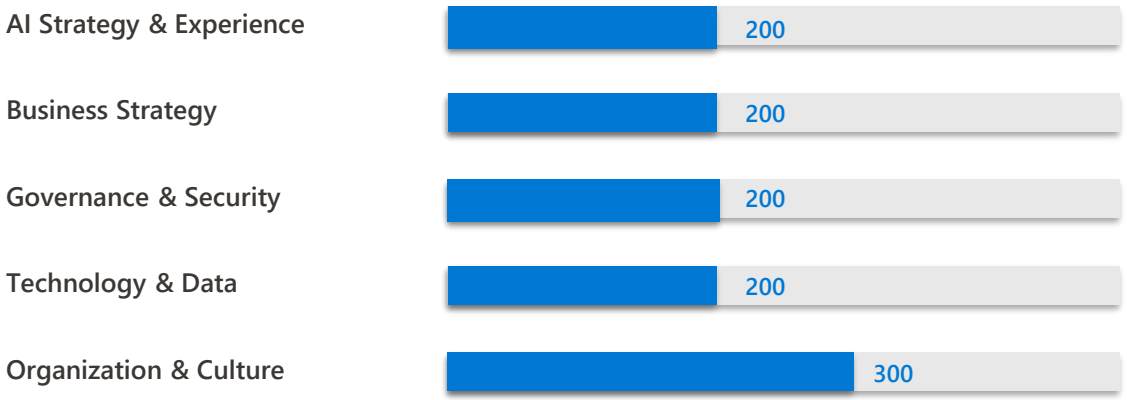

KEY INSIGHT

This pattern needs high culture maturity (300) because you're asking every employee to change how they work. Technology is the easy part — adoption is the challenge. Without leadership role-modelling and continuous enablement, licenses don't become usage.

SCALE-BREAKER

Organization & Culture — if people aren't enabled and encouraged, adoption stalls regardless of technology readiness.

VALUE DELIVERED

- Broad **productivity uplift** and increased confidence with AI in daily work.
- More **focus on high-value work** by reducing context switching and information-seeking.
- **Higher quality** outputs, not just faster task completion.
- Faster, more **confident decisions** through better access to context and insights.
- More **effective collaboration**, focused on decisions not coordination

Employee AI Enablement: What You Need

Practical requirements and operating model guidance

YOU DON'T NEED

- → Process redesign — agents augment, they don't replace workflows
- → Domain ownership models — individuals own their own output
- → Federated governance — centralized policies are sufficient
- → Multi-agent orchestration — these are single-purpose assistants
- → Heavy IT involvement — this is a user-enablement play

YOU DO NEED

- → Leadership role-modelling — leaders must visibly use and champion agents
- → Continuous enablement — training, tips, community, not just license deployment
- → Standardized platforms — consistent tools reduce shadow IT risk
- → Lightweight telemetry — track adoption, usage patterns, and value signals
- → Clear acceptable-use policies — what agents can access and what they can't
- → Change management — people need permission and encouragement to change habits

RECOMMENDED COE STRUCTURE

Centralized

A central CoE sets guardrails, manages platforms, runs community programs, and tracks adoption metrics. Individual employees build and use agents within these guardrails. The CoE's primary job is enablement, not control.

Target maturity: AI Strategy & Experience: 200 • Business Strategy: 200 • Governance & Security: 200 • Technology & Data: 200 • Organization & Culture: 300

Business Expert Empowerment

PATTERN 2

Capture, apply, and scale expert knowledge across the organization — without automating decisions or removing human judgment. The expert's knowledge becomes available at scale, but the expert remains accountable for the quality of that knowledge.

WHAT AGENTS DO

Agents surfaces expert knowledge on demand: answering policy questions, providing recommendations based on domain standards, interpreting guidelines, and flagging exceptions. They act as a scalable proxy for the expert's judgment.

WHAT THE HUMAN DOES

The domain expert defines the rules, curates the knowledge, and validates agent accuracy. They shift from answering every question to teaching the agent and auditing its output. The expert owns the agent's credibility.

EXAMPLE USE CASES

- → Policy and compliance Q&A agents
- → Engineering standards and best-practice guidance
- → Risk assessment decision support
- → Regulatory interpretation assistants
- → Quality criteria and inspection guidance

OPERATING SHIFTS

People

Answering every question → Owning judgment and teaching the agent

Agents

Basic Q&A → Recommendation + escalation to expert

Governance

Source approval → Decision boundaries + knowledge quality controls

Metrics

SME hours spent → Deflection rate + answer accuracy + expert time freed

Business Expert Empowerment: Maturity Profile

Target maturity across five capability drivers

TARGET MATURITY PROFILE

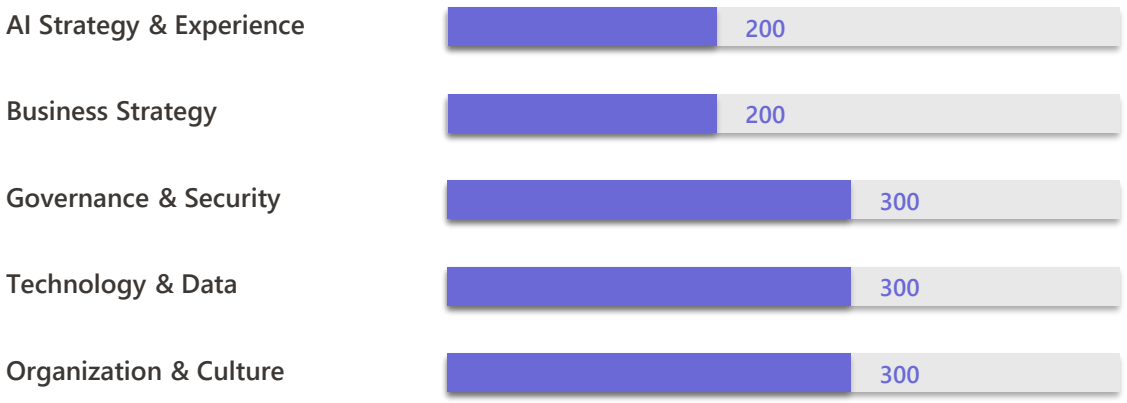

KEY INSIGHT

The agent's credibility IS the product. If the agent gives wrong expert advice, you damage the expert's reputation and potentially the business. Maturity depth concentrates around governance and knowledge quality — not technology.

SCALE-BREAKER

Technology & Data — specifically knowledge quality controls. If you can't guarantee the source documents are authoritative, current, and complete, the agent's output is unreliable.

VALUE DELIVERED

- Expert guidance available **at scale**, on demand.
- **Reduced dependency** on a small number of SMEs.
- Faster, **more consistent decisions** across the organization.
- Preservation of institutional **knowledge** that would otherwise be a bottleneck or risk.
- Experts spend more time on **judgment and exceptions**, not repetitive questions.

Business Expert Empowerment: What You Need

Practical requirements and operating model guidance

YOU DON'T NEED

→ End-to-end process automation — the agent advises, it doesn't execute

→ Process orchestration — this is about knowledge, not workflows

→ Outcome ownership transfer — the expert remains accountable

→ Cross-system integration — the agent works within the expert's domain

→ Enterprise-wide rollout — start with one domain, prove value, expand

YOU DO NEED

→ Authoritative knowledge sources — the agent is only as good as its data

→ Named expert ownership — a specific person, not 'the team', owns accuracy

→ Escalation rules — when the agent doesn't know, it must say so and route to the expert

→ Feedback loops — the expert continuously reviews and improves agent responses

→ Explicit decision boundaries — what the agent can recommend vs. what requires human judgment

→ Knowledge quality monitoring — track accuracy, staleness, and user trust over time

RECOMMENDED COE STRUCTURE

Federated

Experts build and own their agents within their domain. The CoE provides patterns, knowledge guardrails, evaluation standards, and lifecycle discipline. The CoE ensures consistency across domains without centralizing ownership.

Target maturity: AI Strategy & Experience: 200 • Business Strategy: 200 • Governance & Security: 300 • Technology & Data: 300 • Organization & Culture: 300

Workplace & IT Services

PATTERN 3

Agents operate workplace services end-to-end — HR, IT helpdesk, Finance, Facilities — improving reliability, speed, and employee experience. These agents don't just answer questions; they execute service workflows.

WHAT AGENTS DO

Agents handle intake, triage, and routine execution of internal service requests: processing leave requests, provisioning access, validating expenses, answering payroll questions, routing procurement workflows. They operate as a service, not an assistant.

WHAT THE HUMAN DOES

Service owners define service levels, monitor quality, handle exceptions, and manage escalations. They shift from processing every ticket to governing the service and improving it over time.

EXAMPLE USE CASES

→ IT helpdesk — password resets, access provisioning, device troubleshooting

→ HR services — leave requests, onboarding workflows, policy inquiries

→ Finance — expense validation, invoice processing, budget inquiries

→ Facilities — room booking, maintenance requests, workplace services

→ Procurement — purchase order routing, vendor onboarding, approval workflows

→ Data preparation — cleaning, tagging, validating large document sets

OPERATING SHIFTS

People

Task execution → Service ownership and exception handling

Agents

Intake + triage → Routine end-to-end execution

Governance

Tool-level controls → Service-level controls with SLAs

Metrics

Tickets handled → Resolution time + satisfaction + cost per resolution

Workplace & IT Services: Maturity Profile

Target maturity across five capability drivers

TARGET MATURITY PROFILE

KEY INSIGHT

This pattern requires stronger maturity in Business Strategy (400) and Governance (400) than you might expect — because agents are running operational services, not just helping individuals. Service reliability and escalation discipline become critical.

SCALE-BREAKER

Business Strategy — specifically end-to-end service design. If you automate individual tasks without redesigning the service flow, you get islands of automation that don't connect. Design the service, then build the agents.

VALUE DELIVERED

- **Faster resolution** and reduced service backlogs
- More **consistent**, predictable service delivery
- **Lower cost-to-serve** through automation of routine execution
- **Improved employee experience** across HR, IT, and Finance
- Service **teams focus on outcomes**, quality, and improvement — not ticket handling

Workplace & IT Services: What You Need

Practical requirements and operating model guidance

YOU DON'T NEED

→ Domain product ownership — these are shared services, not business products

→ External identity models — these serve internal employees only

→ Full organizational culture transformation — service teams adapt, not the whole org

→ Custom-built AI infrastructure — platform capabilities are sufficient

→ Perfection before launch — start with one service, iterate based on data

YOU DO NEED

→ End-to-end service ownership — someone accountable for the full service lifecycle

→ Decision rights — which requests can the agent resolve autonomously vs. escalate

→ Monitoring and telemetry — uptime, accuracy, resolution time, user satisfaction

→ Run-model with escalation paths — what happens when the agent can't resolve

→ Service-level agreements — define what 'good' looks like and measure against it

→ Integration contracts — clear inputs, outputs, and handoffs between agents and systems

RECOMMENDED COE STRUCTURE

Centralized → Hybrid

Start centralized for control and consistency. As services mature, shift to hybrid: central platform and security guardrails with service owners running day-to-day operations. The CoE owns the platform; service teams own the agents.

Target maturity: AI Strategy & Experience: 400 • Business Strategy: 400 • Governance & Security: 400 • Technology & Data: 300 • Organization & Culture: 300

Core Business Process Transformation

PATTERN 4

Agents run core enterprise processes end-to-end across multiple systems. These are business-critical workflows where agents make decisions, not just suggestions, with direct impact on revenue, cost, and customer experience.

WHAT AGENTS DO

Agents orchestrate complex workflows across systems: processing claims, managing orders, coordinating supply chains, executing financial close processes. They make routine decisions autonomously and escalate exceptions to humans.

WHAT THE HUMAN DOES

Humans shift from doing the work to governing the system. They define autonomy limits, review exception cases, monitor business outcomes, and continuously improve the process. Accountability stays with the business, not IT.

EXAMPLE USE CASES

→ Claims processing — intake, verification, risk scoring, payment

→ Order-to-cash — order validation, fulfillment coordination, invoicing

→ Financial close — reconciliation, accrual calculations, variance flagging

→ Manufacturing quality — inspection triage, defect classification, routing

→ Supply chain — demand sensing, inventory optimization, supplier coordination

→ Customer onboarding — KYC verification, account setup, welcome workflows

OPERATING SHIFTS

People

Doing work → Owning performance and governing the system

Agents

Executing steps → Orchestrating end-to-end across systems

Governance

Tool-use policies → Autonomy limits + decision rights framework

Metrics

Productivity gains → Cycle time + throughput + accuracy + business value

Core Business Process Transformation: Maturity Profile

Target maturity across five capability drivers

TARGET MATURITY PROFILE

KEY INSIGHT

This pattern demands depth everywhere — there's no driver you can skip. 400-500 level maturity across all five drivers. This is where agents stop being tools and start being part of the business. It directly impacts profit and loss and requires high organizational maturity.

SCALE-BREAKER

Everything — but if forced to choose one: Business Strategy. Without formal process redesign and KPI-linked orchestration, you're automating broken processes faster.

VALUE DELIVERED

- **Step-change improvements** in cycle time, throughput, and quality.
- **Lower operating cost** through reduced handoffs and waiting.
- More resilient, predictable core operations.
- Direct impact on **business KPIs** (revenue, cost, service levels).
- Teams shift from doing work to **owning performance.**
- **Federated innovation embedded** within business teams.

Core Business Process Transformation: What You Need

Practical requirements and operating model guidance

YOU DON'T NEED

→ Individual enablement programs — this isn't about personal productivity

→ Lightweight governance — this pattern has profit and loss impact and needs formal controls

→ Centralized IT delivery — the business must own the outcomes

→ Incremental automation — you're redesigning the process, not patching it

→ Skipping change management — people's roles fundamentally change

YOU DO NEED

→ Process redesign with human accountability — who owns the outcome, not just the agent

→ Decision rights framework — which decisions can agents make autonomously vs. require approval

→ Autonomy limits — clear boundaries for agent actions, especially financial and customer-impacting

→ Production-grade monitoring — SLAs, accuracy tracking, drift detection, not just uptime

→ Federated governance — business units own agents, CoE provides standards and oversight

→ Change management — the humans in this process need a new role, not just a new tool

RECOMMENDED COE STRUCTURE

Federated

Business units own agent delivery and outcomes. The CoE provides standards, enablement, and governance-by-exception. This requires mature teams in the business who can own agent lifecycle — not just build agents, but run, monitor, and improve them.

Target maturity: AI Strategy & Experience: 500 • Business Strategy: 500 • Governance & Security: 400 • Technology & Data: 400 • Organization & Culture: 400

External Engagement

PATTERN 5

Agents interact directly with customers, partners, or ecosystem stakeholders — crossing the enterprise trust boundary. Every interaction affects brand, reputation, and customer trust. Errors are visible externally.

WHAT AGENTS DO

Agents handle customer-facing interactions: answering support questions, guiding purchases, onboarding partners, providing status updates, resolving routine issues. They represent the organization to the outside world.

WHAT THE HUMAN DOES

Humans define brand tone, set escalation rules, monitor interaction quality, and handle complex or sensitive cases. They shift from delivering every service interaction to overseeing trust and quality at scale.

EXAMPLE USE CASES

→ Customer support — tier-1 inquiry resolution, FAQ, troubleshooting

→ Digital concierge — product guidance, recommendation, onboarding

→ Partner enablement — self-service portals, documentation access

→ Vendor onboarding — compliance verification, document collection

→ Status tracking — order status, delivery updates, service requests

→ Feedback collection — surveys, sentiment analysis, issue escalation

OPERATING SHIFTS

People

Service delivery → Trust oversight and quality governance

Agents

Internal tools → External-facing execution representing the brand

Governance

Usage policies → Disclosure, consent, and identity isolation requirements

Metrics

Engagement volume → Customer satisfaction + trust + resolution quality

External Engagement: Maturity Profile

Target maturity across five capability drivers

TARGET MATURITY PROFILE

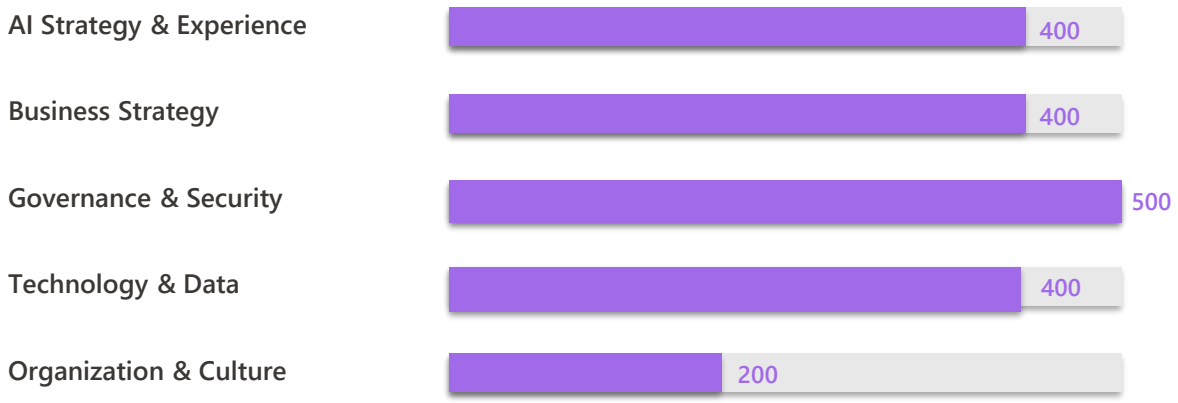

KEY INSIGHT

External engagement requires deeper governance and security maturity (500) than any internal pattern because errors directly impact customers, brand, and potentially regulatory compliance. The trust boundary is real.

SCALE-BREAKER

Governance & Security — specifically identity isolation, disclosure requirements, and interaction monitoring. One bad customer interaction from an unsupervised agent is a brand crisis.

VALUE DELIVERED

- **Differentiated** customer and partner experiences.
- Increased reach and availability without linear cost growth.
- Faster service, higher satisfaction, and **improved engagement**.
- **New digital service** and ecosystem-driven business models.
- **Growth** with resilience and trust at scale

External Engagement: What You Need

Practical requirements and operating model guidance

YOU DON'T NEED

→ BU-led autonomous delivery — external agents need central oversight

→ Local policy variations — brand consistency is non-negotiable

→ Lightweight disclosure — customers must know they're interacting with an agent

→ Skipping identity isolation — external and internal agent access must be separated

→ Deploying without human handoff — every external agent needs an escalation path

YOU DO NEED

→ Identity isolation — strict separation between internal and external agent access

→ Real-time monitoring and abuse detection — anomalous interaction patterns caught immediately

→ Human handoff capabilities — seamless escalation from agent to human for complex issues

→ Consent and disclosure frameworks — transparency about AI use in every interaction

→ Brand-aligned tone and behavior rules — the agent represents your organization

→ Incident response plan — when an external agent gives wrong information, what happens in 15 minutes?

RECOMMENDED COE STRUCTURE

Centralized

Crossing the enterprise boundary requires the tightest controls. A centralized CoE sets decision rights, manages disclosure requirements, monitors interactions, and maintains consistent evidence standards across all external-facing agents.

Target maturity: AI Strategy & Experience: 400 • Business Strategy: 400 • Governance & Security: 500 • Technology & Data: 400 • Organization & Culture: 200

AI-First Capabilities

PATTERN 6

Design net-new capabilities with agents as the core building blocks — not retrofits to existing workflows. These are things that weren't possible before AI. Capabilities improve over time through feedback, learning, and continuous optimization.

WHAT AGENTS DO

Agents operate in sense-decide-act loops: continuously monitoring signals, making autonomous decisions within boundaries, executing actions, and learning from outcomes. These aren't assistants — they're intelligent systems that create new business value.

WHAT THE HUMAN DOES

Humans act as product owners: defining objectives, setting boundaries, monitoring outcomes, and making strategic decisions about capability evolution. They don't do the work — they govern the system that does.

EXAMPLE USE CASES

→ Continuous optimization engines — pricing, inventory, scheduling

→ AI-native decision loops — fraud detection, anomaly response, adaptive routing

→ Market sensing platforms — competitive monitoring, trend analysis, opportunity identification

→ Predictive planning systems — demand forecasting, capacity planning, risk modeling

→ Autonomous workflow generation — agents that design and optimize their own processes

→ Multi-agent orchestration — coordinated teams of specialized agents solving complex problems

OPERATING SHIFTS

People

Operators executing tasks → Product owners governing capabilities

Agents

Executing predefined steps → Sense-decide-act autonomous loops

Governance

Compliance checks → Lifecycle-embedded controls with continuous monitoring

Metrics

Throughput and efficiency → Strategic optionality + learning rate + adaptation speed

AI-First Capabilities: Maturity Profile

Target maturity across five capability drivers

TARGET MATURITY PROFILE

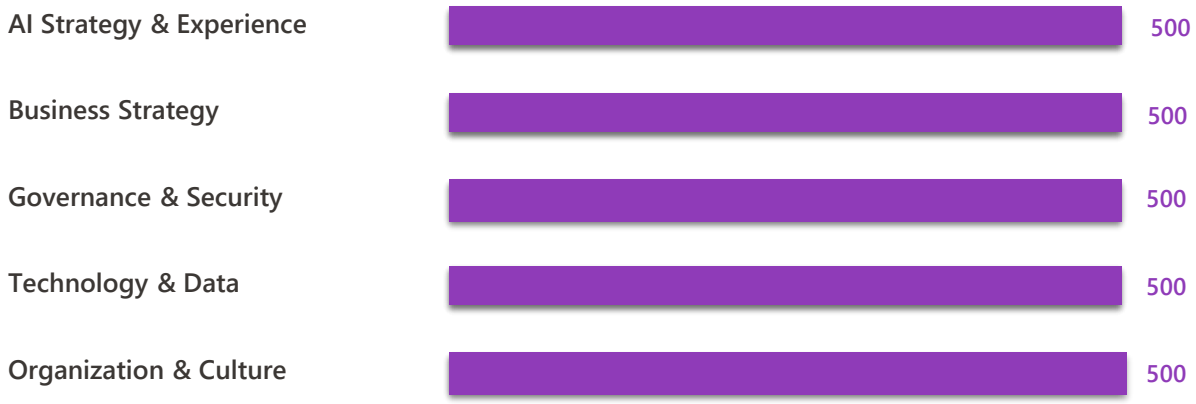

KEY INSIGHT

This pattern demands the highest maturity across all capabilities. There's no existing process to compare against, no baseline to measure improvement from, and no incumbent workflow to guide design. Everything must be built — including how you measure success.

SCALE-BREAKER

Technology & Data — specifically multi-agent orchestration, real-time telemetry, and learning infrastructure. Without a robust technical foundation, autonomous capabilities become unpredictable.

VALUE DELIVERED

- Entirely new sources of **competitive advantage**.
- **Faster adaptation** to market and operational change.
- **Higher strategic optionality** — the business can respond in ways it couldn't before.
- **Differentiation** that competitors cannot easily replicate.
- **Compounding value** as capabilities learn and improve over time

AI-First Capabilities: What You Need

Practical requirements and operating model guidance

YOU DON'T NEED

→ Incremental automation — this isn't about making existing processes 10% faster

→ Central IT delivery — capabilities need dedicated product teams

→ Retrofitting existing workflows — design from scratch for AI-native execution

→ One-time builds — these capabilities must continuously learn and improve

→ Traditional project management — use product management with experiment cycles

YOU DO NEED

→ Product ownership with dedicated teams — each capability needs a product owner and team

→ Autonomy boundaries with continuous monitoring — clear limits that adapt as trust builds

→ Continuous learning loops — the capability improves through every interaction

→ Embedded responsible AI practices — ethics and safety built into the architecture, not bolted on

→ Experimentation culture with fail-safe controls — permission to try, safety nets when things go wrong

→ Multi-agent architecture — design for agent-to-agent coordination from the start

RECOMMENDED COE STRUCTURE

Federated

AI-first capabilities are owned as products, span multiple domains, and require distributed product teams operating within shared standards. The CoE provides architecture patterns, responsible AI frameworks, and orchestration standards — not central delivery.

Target maturity: AI Strategy & Experience: 500 • Business Strategy: 500 • Governance & Security: 500 • Technology & Data: 500 • Organization & Culture: 500

Section 2
The Maturity Model

A diagnostic to find what breaks scale first — not a scorecard

Understanding the Maturity Model

How to use this diagnostic to find your scale-breaker and prioritize action

The Agentic AI Maturity Model assesses your organization's readiness to build, govern, and operate agents at scale. It looks at five capability drivers — each of which can independently block your ability to scale.

The model has five levels (100–500), but the goal is NOT to reach 500 everywhere. Different adoption patterns require different maturity depths across different drivers. Your job is to:

1. Identify which pattern you're pursuing
2. Look up the target maturity profile for that pattern
3. Assess where you are today across the five drivers
4. Find the biggest gap — that's your scale-breaker
5. Focus investment on closing that gap first

The maturity model is a diagnostic, not a scorecard. You're not trying to get a high score — you're trying to find the ONE thing that will break scale first.

The five capability drivers:

AI Strategy & Experience	How deliberately you plan, invest in, and evolve AI across the organization
Business Strategy	How deeply AI is integrated into business processes and outcome measurement
AI Governance & Security	How well you manage risk, compliance, monitoring, and responsible AI
Technology & Data	How mature your platforms, architecture, data quality, and telemetry are
Organization & Culture	How effectively you enable adoption, build skills, and foster AI-positive culture

Learn more: https://aka.ms/AgentMaturityModel

Five Maturity Levels

Each level represents increasing depth of organizational capability

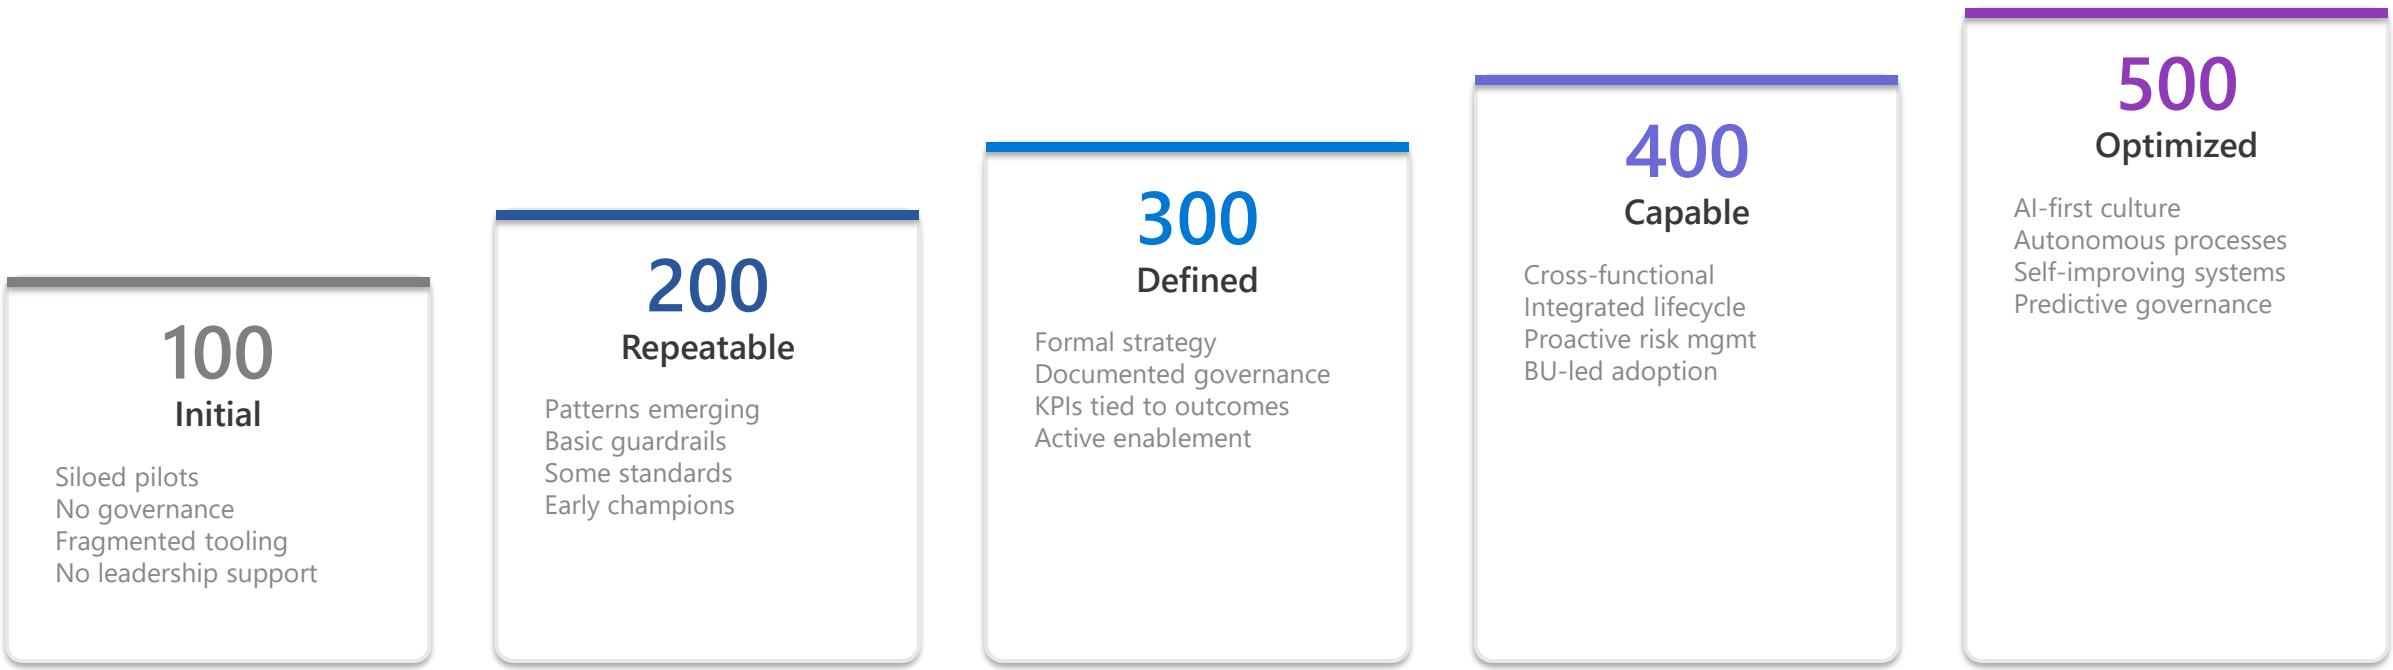

Most organizations are a patchwork — maybe 100 in Business Strategy but 300 in Governance & Security. That unevenness is exactly where scale breaks. Your weakest driver becomes your ceiling, regardless of how strong the others are.

Don't try to move from 200 to 300 everywhere. Find the ONE driver that will break scale first. Fix that. Then reassess.

Each Pattern Demands Different Maturity

The drivers don't change — the required depth does

	AI Strategy & Experience	Business Strategy	Governance & Security	Technology & Data	Organization & Culture
Employee AI Enablement	200	200–300	200	200–300	300
Business Expert Empowerment	200	200	300	300	300
Workplace & IT Services	300-400	400	400	300	200–300
Core Business Transformation	500	500	400	400	400
External Engagement	300–400	400	400–500	400	200
AI-First Capabilities	500	500	500	500	500

Maturity: 200 — Repeatable | 300 — Defined | 400 — Capable | 500 — Optimized

Your pattern determines WHERE you invest — not just how much

AI Strategy and experience

From early experimentation to an agent-first strategy that's embedded in enterprise planning

Where is your organization?

1 Initial

- No AI/agent strategy; experimentation is siloed and disconnected from business goals.
- No executive sponsorship; success is anecdotal and user experience is unplanned.
- Responsible AI awareness is minimal and not integrated into how teams plan or decide.

Signals you'll recognize:

- Different teams try different things; users struggle to know what's trustworthy or "official".

2 Repeatable

- An AI vision is forming; but it's informal and inconsistently communicated.
- Teams are gaining experience; basic agent experiences exist, but usability and consistency vary.
- Responsible AI is recognised, but not yet embedded as a strategic pillar.

Signals you'll recognize:

- Early wins… but the organization can't repeat them consistently beyond a few motivated teams.

3 Defined

- A documented AI/agent strategy aligned to business objectives is in place, with clear sponsorship.
- Objectives and reporting exist; users can discover and use agents more consistently across the org.
- Responsible AI is integrated into strategic planning and decision-making.

Signals you'll recognize:

- Teams share a common direction; decisions are increasingly made using shared principles and measures.

4 Capable

- AI strategy is embedded into enterprise planning, with cross-functional alignment and aligned sub-plans.
- User experiences become more seamless and contextual, enabling broader adoption and satisfaction.
- Responsible AI is positioned as a strategic advantage and is integrated into communications and oversight.

Signals you'll recognize:

- Business units execute within a shared strategy; scale decisions are made using outcome signals and governance.

5 Efficient

- The organization operates with a living, adaptive AI strategy and executive accountability for outcomes.
- AI experiences continuously improve based on telemetry and feedback, agents feel like "natural" collaborators.
- Responsible AI is deeply embedded in culture and treated as a core business value.

Signals you'll recognize:

- Strategy updates are routine, driven by real outcomes; the organization leads with trust and clarity as capability grows.

Key enablers

- Strategy on a page
- Experience principles + consistency
- Responsible AI embedded by design
- Feedback and measurement loops

Business strategy

From task-level pilots to agent-orchestrated processes with measurable, repeatable business value.

Where is your organization?

1 Initial

- Manual, human-led processes; no end-to-end redesign for agent collaboration.
- If used, agents assist but don't orchestrate work or trigger systems in a coordinated way.
- No formal outcomes tracking success is anecdotal and baselines/success criteria are missing.

Signals you'll recognize:

- Lots of demos, limited sustained impact.

2 Repeatable

- Pilots improve single steps; incremental improvements without coordinated end-to-end transformation.
- Some metrics exist, but measurement is inconsistent and not standardized.
- Value often assessed after delivery, not planned upfront or tied to enterprise OKRs/KPIs.

Signals you'll recognize:

- Teams can repeat a few wins, but each new initiative reinvents measurement.

3 Defined

- Workflows are redesigned end-to-end so agents participate in execution, with humans focusing on judgement and exceptions.
- Human-agent roles and boundaries are documented for priority processes.
- Agents have defined KPIs (time saved, error reduction, satisfaction) and value is tracked per project.

Signals you'll recognize:

- You can point to "how work runs now", and how success is measured.

4 Capable

- Agents orchestrate multi-step, cross-system workflows; human-led, agent-operated processes become standard in key functions.
- Regular value reporting to leadership; proven ROI across multiple agents.
- Metrics include operational and strategic value, and feedback informs improvement.

Signals you'll recognize:

- Scaling decisions are portfolio-based, not enthusiasm-based.

5 Efficient

- Core business processes are agent-operated, adaptive, and continuously optimized; higher autonomy with human oversight.
- Real-time enterprise view of AI value; decisions to scale/modify/retire are data driven.
- Value metrics span outcomes, experience, and trust/risk indicators.

Signals you'll recognize:

- AI value reporting is routine and drives strategy, not just status updates.

Key enablers

- AI-first process design
- Decision rights + orchestration patterns
- Value measurement discipline
- Make value visible

AI governance and security

From ad hoc guardrails to tiered, automated governance and predictable operations as agents scale.

Where is your organization?

1 Initial

- No AI-specific standards; agents run without formal oversight, risk assessment, or compliance checks; ownership/decision rights are unclear.
- No formal support model; little/no monitoring, improvement loop, or structured incident response.
- No formal Responsible AI awareness or practices

Signals you'll recognize:

- Agent creation and sharing happens without a consistent approval path, and ownership/decision rights aren't clear

2 Repeatable

- Basic controls/policies exist but are inconsistently applied; reviews are manual and reactive; early environment separation may exist.
- Basic monitoring exists; support is reactive and depends on a few individuals; runbooks may be informal.
- Early checklists/reviews appear but practices are inconsistent.

Signals you'll recognize:

- Some teams use dev/test/prod separation and reviews, while others don't — governance depends on individual diligence

3 Defined

- Practices are documented and enforced; agents are classified by purpose/criticality/autonomy.
- Formal operations model exists; monitoring and incident management are defined; continuous improvement loops start emerging.
- Higher-risk agents require stronger Responsible AI scrutiny

Signals you'll recognize:

- You can point to a defined classification/tiering approach that changes controls based on risk and criticality.

4 Capable

- Governance is risk-based and partially automated; standards are central, approvals can be delegated for low-risk agents within guardrails.
- Operations become proactive (alerts, anomaly detection); monitoring and optimisation are part of the operating rhythm.
- Responsible AI is embedded into lifecycle gates and day-to-day practice.

Signals you'll recognize:

- Low-risk agents move faster because approvals are delegated within guardrails, while critical agents follow enterprise ALM/security rigor.

5 Efficient

- Agents are treated as tiered digital services with differentiated SLAs/controls; governance continuously adapts based on usage, risk, and regulation.
- Predictive, self-optimising operations.
- Responsible AI is internalised culturally and operationally; trust and ethics are part of strategic performance discussions

Signals you'll recognize:

- Continuous compliance and "always-on" improvement loops are standard, not special cases.

Key enablers

- Risk tiering + classification
- Inventory + ownership
- Enforceable guardrails
- Observability + incident readiness

Technology and data

Redesign processes for human–agent collaboration and make value measurable and repeatable.

Where is your organization?

1 Initial

- Agent work is exploratory and fragmented; teams experiment without a defined technology plan.
- Data access is unplanned, often limited to Microsoft 365 documents or direct system calls with no consistent retrieval strategy.
- No consistent platform, ALM, or integration standards; solutions are fragile and hard to scale.

Signals you'll recognize:

- Different teams are experimenting in different ways, and there's no shared standard architecture across agents.

2 Repeatable

- Teams converge on a small set of platforms, but choices are inconsistent and driven by team preference vs use-case fit.
- Microsoft 365 content is accessible, but structured business data remains siloed with limited approved connectors.
- Some environment separation exists, but ALM, telemetry, and documentation are inconsistent.

Signals you'll recognize:

- Some teams can build agents faster than others because practices vary by team, not by an enterprise standard

3 Defined

- A documented technology plan exists; the org consistently distinguishes between SaaS agents, Copilot Studio agents, and advanced build paths.
- Data architecture follows a clear pattern.
- Standard platforms, reference architectures, integration patterns, and ALM are defined and reused; teams use a structured design framework.

Signals you'll recognize:

- Production agents follow defined ALM practices rather than ad hoc promotion.

4 Capable

- Foundations are enterprise-grade with clear visibility into systems/APIs/data used across workflows.
- Secure-by-design access patterns, centralised telemetry, and automated evaluations are standard; deployments are automated and reliable.
- Central monitoring provides visibility into agent behaviour and performance across the organisation.

Signals you'll recognize:

- You have central monitoring/telemetry that shows how agents behave and perform across teams.

5 Efficient

- Technology and data foundations continuously evolve based on telemetry and emerging agent patterns.
- Patterns are maintained as shared enterprise assets; federated teams build independently while guardrails enforce quality by default.
- Architecture supports complex, multi-agent scenarios and federated development at scale.

Signals you'll recognize:

- Teams can deliver independently because guardrails enforce quality by default, not through central bottlenecks

Key enablers

- Standard platforms + reference architectures
- ALM + automated release discipline
- Governed data & integration access
- Telemetry, observability & evaluation

Organization and culture

Redesign processes for human–agent collaboration and make value measurable and repeatable.

Where is your organization?

1 Initial

- AI adoption exists only in isolated experiments or pilots.
- No shared enterprise narrative for AI/agents; leadership sponsorship is weak or absent.
- Agents are seen as optional/technical, with no clear ownership for adoption, value, or risk; learning is informal and self-directed.

Signals you'll recognize:

- People try agents in pockets, no shared expectation for how AI should be used day-to-day.
- Upskilling happens through self-directed learning rather than a structured programme.

2 Repeatable

- Interest is growing, but adoption depends on motivated individuals or teams.
- Training and communities exist sporadically, without a structured programme or operating model.

Signals you'll recognize:

- Some groups have active champions or informal communities, but coverage is uneven across the organization.
- Teams interpret "what's allowed" differently because ownership and decision rights aren't explicit.

3 Defined

- A central team or Center of Excellence (CoE) provides standards and enablement while execution is federated.
- Structured onboarding, learning paths, community events, and knowledge repositories/showcases are in place; leadership endorsement is present but uneven.

Signals you'll recognize:

- New users/makers have a clear onboarding path.
- Communities and showcases exist as a repeatable rhythm, not ad hoc events.

4 Capable

- Agent-assisted work is standard practice across functions; incentives/expectations reinforce responsible use.
- Culture supports experimentation within clear guardrails; business units proactively propose agent-enabled improvements.
- Communities are active and self-driven, with regular showcases and hackathons that emphasise responsible practice.

Signals you'll recognize:

- You see regular, self-sustaining community activity (showcases/hackathons) rather than programme-led pushes.

5 Efficient

- The organisation operates as an agent-first enterprise; grassroots ideas are rapidly surfaced, governed, and scaled.
- Culture, leadership, incentives, and learning are fully aligned around responsible AI practices; communities reinforce standards and innovation.
- Employees demonstrate mature Responsible AI habits and ethical reasoning.

Signals you'll recognize:

- Innovation scales broadly because the organisation can surface and govern ideas quickly without losing trust.

Key enablers

- Visible executive sponsorship + shared narrative
- Baseline operating model + decision rights
- Structured enablement + community nurture
- Adoption signals + reinforcement loops

The Top 5 Scale-Breakers

Common signals that indicate where your maturity gaps are blocking scale

These are the five most common patterns we see across organizations trying to scale agents. Most organizations recognize at least three. The question is: which one will you fix first?

Signal	Root Cause	Fix
Many pilots, no portfolio	No outcome-led portfolio (agents aren't tied to measurable business outcomes)	**Pick 1–2 outcomes + 1–2 patterns, name business owners, and define success metrics (what improves, by when)**
One-off agents, no reuse	Weak reusable foundations (data quality/ownership + integration standards + telemetry foundations aren't consistent)	**Standardise reference architecture + integration approach for the chosen pattern; establish telemetry baseline and a "known-good" knowledge source strategy.**
Great demos, low adoption	The AI experience isn't designed end-to-end (users don't know when to use it, what it can do, or how to validate results).	**Define golden paths for the top scenarios (how users engage, what's automated vs human-approved, and how exceptions are handled).**
Licenses ≠ usage	Enablement and change motion isn't systematic (no community, training, champions, or incentives tied to new ways of working).	**Launch a nurture motion: role-based enablement + community cadence + 'safe path is the easy path' templates.**
Shadow agents appearing	Governance is not operational (decision rights, escalation, and evidence aren't defined per risk tier).	**Implement a minimum baseline: named owner + audit trail + release gate + monitoring + escalation (risk-tiered).**

Section 3
The Agentic Center of Excellence

The execution engine for operating and scaling agents

What is a Center of Excellence?

Understanding the CoE's role in operationalizing agents at scale

A Center of Excellence (CoE) is the organizational vehicle that turns strategy and intent into repeatable, trusted execution. It's not a committee, not a silo, and not just governance. It's the team — and the operating rhythm — that makes it possible for the rest of the organization to build, deploy, and run agents safely.

Without a CoE, each team builds agents independently with different standards, different quality levels, and no shared learning. The first 5 agents work fine. By agent 50, you have a governance crisis. The CoE prevents this by providing four functions: Govern, Enable, Optimize, and Scale.

A CoE turns "we built an agent" into "we run agents as a business capability"

A CoE is NOT:

- ❌ The Gatekeeper — Reviews everything, enables nothing, becomes the bottleneck. Teams route around it.
- ❌ The Ghost — Launched with fanfare, no operating rhythm, no decision authority.
- ❌ The Perfectionist — Won't let anyone ship until governance is 'complete.' Governance is never complete.

What a CoE actually does

Four functions that turn intent into repeatable execution

Governs

Guardrails, not gates

Risk-proportionate controls that enable speed, not block it. Release gates ensure nothing goes to production without review. Audit logs track who built what, who approved it, and what it does. Compliance is continuous, not a one-time checkbox.

Release gates • Audit trails • Compliance monitoring • Risk classification

Enables

Golden paths for makers

Make the right way the easy way. Provide templates, training, community, and best practices so teams build well the first time. A maker shouldn't have to figure out governance on their own — the CoE should make safe building the path of least resistance.

Templates • Training • Community • Best practices • Office hours

Optimizes

Lifecycle management

Agents are products, not projects. They need ongoing monitoring, evaluation, and improvement. The CoE ensures every agent has health checks, accuracy tracking, and a plan for what happens when quality drifts.

Monitoring • Evaluation • Accuracy tracking • Drift detection • Improvement cycles

Scales

Repeatable execution

Turn heroic one-offs into a repeatable pipeline. Intake → Triage → Build → Deploy → Run → Improve → Retire. The fiftieth agent should be easier to build, deploy, and operate than the fifth. Patterns, reuse, and standardization make this possible.

Intake pipeline • Patterns • Reuse • Standardized architecture • Portfolio management

An agent is a product, not a project

Agents need a lifecycle

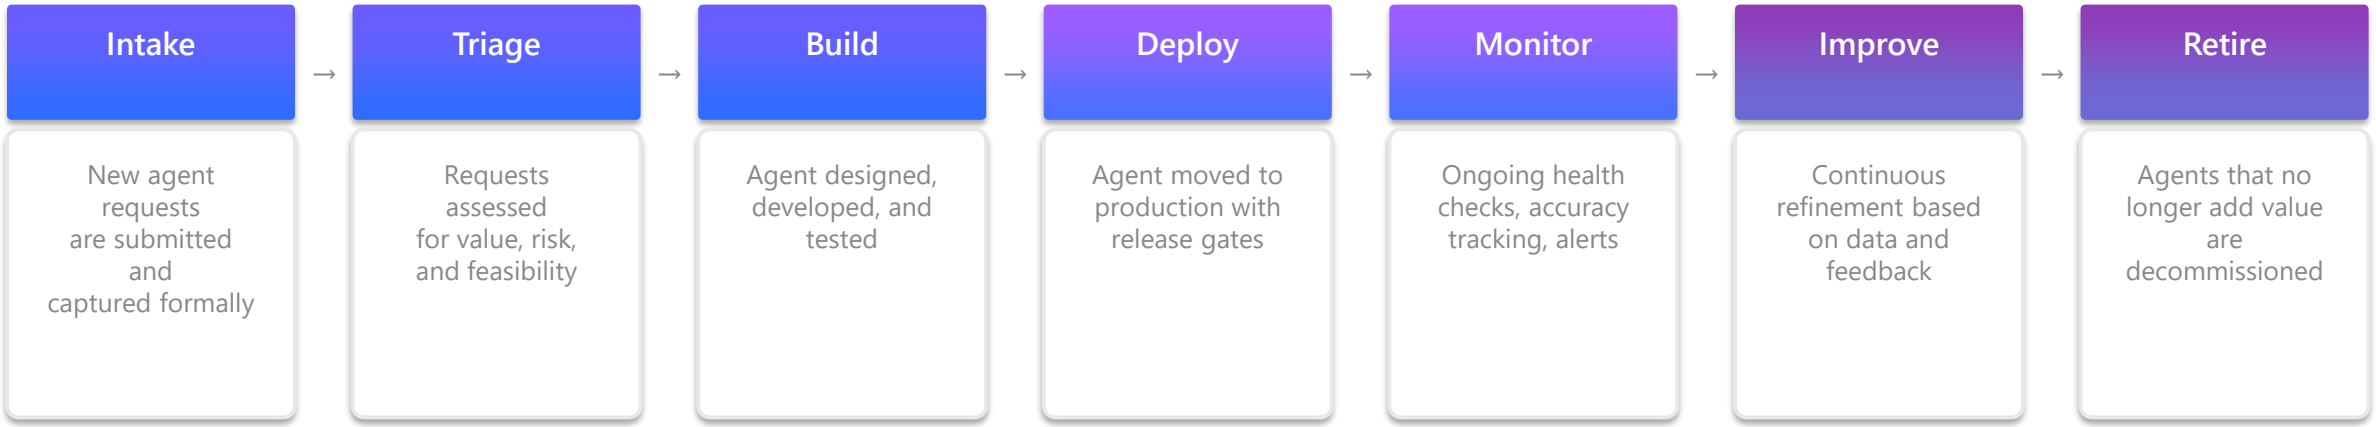

Every agent in production without monitoring and an improvement plan is accumulating risk in your tenant. Knowledge goes stale. Source documents get updated in a place the agent does not have access to. User patterns change. Integrations break. Without lifecycle management, agents don't fail dramatically — they slowly drift, giving increasingly wrong answers with full confidence. That's worse than a crash, because nobody notices.

The CoE ensures every agent has an owner, a monitoring plan, and a defined path to improvement or retirement.

Match your CoE structure to your adoption pattern

The right structure depends on your intent, not your org chart

Centralized

Control

A single team sets rules, delivers, monitors. Best when you're early, when risk is high, or when agents cross enterprise boundaries. Provides maximum consistency and control.

BEST FOR:

→ Employee AI Enablement

→ External Engagement

→ Workplace & IT Services

IN PRACTICE:

In practice: governance and most delivery sit in one CoE team for control and consistency. Often best early in the journey or in regulated industries.

Hybrid

Balance

Central team sets standards and supplies expertise; local teams build within guardrails. A matrix model where the CoE enables rather than controls.

BEST FOR:

→ Business Expert Empowerment

→ Workplace & IT Services

IN PRACTICE:

In practice: the CoE defines patterns, templates, and quality standards. Domain experts and makers build agents with CoE support. Knowledge quality remains centrally governed.

Federated

Scale

Business units own agent outcomes and delivery. The CoE provides standards, enablement, and oversight. Requires mature teams who can own agent lifecycle independently.

BEST FOR:

→ Core Business Processes

→ AI-First Capabilities

IN PRACTICE:

In practice: BU teams own agents end-to-end. The CoE provides architecture standards, responsible AI frameworks, and governs by exception. Scale comes from distributed ownership.

CoE Roles & Responsibilities

Six roles that make an Agentic CoE work — no single role can scale agents alone

Executive Sponsor

Strategic

Sets ambition and guardrails. Aligns funding and priorities. Removes organizational blockers. Champions AI adoption visibly.

Business Owner

Value

Owns value, adoption, and KPIs per domain. Decides: scale, stop, or iterate. Ensures agents land in real workflows — not just demos.

CoE Lead / AI Program Mgr

Operating

Runs intake, prioritization, and operating rhythms. Single point of accountability for 'how scale works.' Connects strategy to execution.

Agent Product Owner

Product

Translates strategy into an agent portfolio. Prioritizes scenarios. Owns the roadmap. Connects outcomes to delivery.

Platform & Operations

Run

Manages environments, monitoring, support, and incident response. Ensures reliability. Owns the technical foundation agents run on.

Security / Risk / Compliance

Trust

Owns risk tiering, responsible AI, and high-risk decision reviews. Ensures governance evolves with the agent portfolio, not after it.

Start with people covering these roles part-time. Formalize as the agent portfolio grows.

Domain-Level Roles

Who owns what at the edge — makers, experts, and service owners

Makers

Build & iterate agents in real workflows. Use Copilot Studio and Agent Builder in Microsoft 365 Copilot to create agents within CoE guardrails. Test, refine, and improve based on user feedback.

ACCOUNTABLE FOR:

Build within guardrails

Domain Experts

Define the rules, criteria, and knowledge that agents use. Validate agent accuracy. Own the credibility of agent output. Shift from doing repetitive work to teaching agents and auditing exceptions.

ACCOUNTABLE FOR:

Own domain knowledge quality

Service Owners

Own end-to-end service delivery for agent-powered services. Define SLAs. Manage escalation paths. Handle incidents. Ensure operational reliability and continuous improvement.

ACCOUNTABLE FOR:

Own service reliability

Adoption Leads

Drive skills and behavior change across the organization. Run enablement programs, champions networks, and community events. Track adoption metrics and identify barriers.

ACCOUNTABLE FOR:

Drive behavior change

The CoE centralizes HOW scale works — standards, telemetry, escalation.
It does NOT centralize WHO builds everything. Makers, experts, and service owners operate at the edge within the guardrails the CoE provides.

RACI: Who Does What Across the Agent Lifecycle

Example

R = Responsible (does the work) • A = Accountable (owns the outcome) • C = Consulted • I = Informed

	Exec Sponsor	Business Owner	CoE Lead	Agent Product Owner	Platform & Ops	Security & Risk	Makers / Experts
Agent strategy & prioritization	A	R	C	R	I	C	I
Use case intake & triage	I	C	A	R	C	C	R
Agent design & build	I	C	C	A	I	C	R
Knowledge curation & quality	I	A	C	C	I	C	R
Release gate / go-no-go	I	C	A	C	R	R	I
Production deployment	I	I	C	C	R	C	I
Ongoing monitoring	I	I	A	C	R	C	I
Incident response	I	C	A	C	R	R	I
Performance review & improvement	I	A	C	R	C	C	R
Agent retirement	I	A	R	C	R	C	C

R — Responsible — does the work

A — Accountable — owns the outcome

C — Consulted — provides input

I — Informed — kept in the loop

How roles shift across adoption patterns

The roles don’t change. The weight shifts. Same title, different job.

	Employee Enablement	Expert Empowerment	Workplace & IT Services	Core Business Process	External Engagement	AI-First Capabilities
Executive Sponsor	Sets tone, removes blockers	Funds SME time	Owns service transformation	P&L accountability	Brand & legal risk owner	New capability bet sponsor
Business Owner	N/A (individual users)	Domain expert themselves	Service delivery lead (HR/IT)	Process owner (e.g. VP Supply Chain)	Customer experience lead	Product GM / innovation lead
Agent Product Owner	CoE manages centrally	Expert + CoE partner	Service team owns, CoE supports	BU owns, CoE governs by exception	CoE owns centrally	Dedicated product team
Platform & Ops	Standard shared platform	Standard + knowledge infrastructure	SLA monitoring, incident mgmt	Cross-system orchestration	Uptime, compliance logging	Custom bespoke architecture
Security / Risk	Lightweight data policies, data boundaries	Knowledge accuracy, escalation	PII handling, access control	Decision audit trails, rollback	Disclosure, consent, evidence	Autonomy bounds, RAI
Adoption Lead	Primary focus: behavior change	Expert community building	Service rollout comms	Process change management	Customer comms strategy	Stakeholder alignment

Light involvement Moderate involvement Heavy involvement

Every pattern has a make-or-break role

Get this one role wrong, and the pattern stalls — regardless of the technology.

Employee Enablement

Adoption Lead

It's a behavior change problem, not a tech problem. If people don't change habits, nothing scales.

Expert Empowerment

Domain Expert

If the expert disengages, the agent loses credibility. Their reputation is on the line.

Workplace & IT Services

Service Owner

Must shift from processing tickets to governing a service. The operating model changes completely.

Core Business Process

Business Owner

Accountability must sit in the business, not IT. This is a P&L decision, not a tech decision.

External Engagement

Security / Risk

One bad customer interaction is a brand incident. Governance can't be an afterthought.

AI-First Capabilities

Agent Product Owner

This is a product, not a project. It needs product management discipline and a roadmap.

Decision Rights: Central vs. Delegated

Centralize HOW scale works — not WHO builds everything

CENTRALIZED — CoE OWNS

Decisions that must be consistent across the organization

- → Platform and environment strategy
- → Security and compliance policies
- → Architecture standards and reference patterns
- → Release readiness / go-no-go gate criteria
- → Monitoring standards and alerting thresholds
- → Agent risk classification tiers
- → Autonomy limits — what agents can/can't do autonomously
- → Responsible AI guidelines and review requirements

DELEGATED — DOMAIN OWNS

Decisions that require domain context and proximity to the work

- → Domain prioritization — which use cases matter most
- → Agent design decisions within architecture standards
- → Knowledge curation — what content the agent uses
- → Day-to-day operations for lower-risk agents
- → User experience and interaction design
- → Domain-specific success metrics and KPIs
- → Continuous improvement — what to adjust and when
- → Delivery execution — building and iterating the agent

The principle: centralize standards and decision rights so every team benefits from shared guardrails.
Delegate execution and domain decisions so teams closest to the work can move fast within those guardrails.
The CoE's job is to make the safe path the easy path — not to own every build.

Risk-Tiered Governance

Not every agent needs the same level of oversight — match controls to risk

Tier 1 — Low Risk

Individual productivity agents
(drafting, summarization, research)

REQUIRED CONTROLS:

- → Named owner
- → Basic monitoring
- → Standard release checklist
- → Self-service deployment within guardrails

APPLIES TO:

Employee AI Enablement

Tier 2 — Medium Risk

Expert knowledge agents and
internal service agents

REQUIRED CONTROLS:

- → Named owner + domain expert validator
- → Knowledge quality monitoring
- → Formal release gate with review
- → Accuracy tracking and feedback loops

APPLIES TO:

Expert Empowerment
Workplace & IT Services

Tier 3 — High Risk

Business-critical and
external-facing agents

REQUIRED CONTROLS:

- → Named owner + formal process owner
- → Production-grade SLA monitoring
- → Security review + RAI assessment
- → Decision rights framework
- → Incident response plan
- → Quarterly maturity review

APPLIES TO:

Core Business Processes
External Engagement
AI-First Capabilities

Governance should be proportionate to risk. Over-governing low-risk agents slows adoption. Under-governing high-risk agents creates liability. Match the tier to the pattern.

Integrating with Existing Governance

The Agentic CoE connects to — not replaces — your existing governance structures

Most organizations already have governance structures for IT, security, data, and low-code platforms. The Agentic CoE should integrate with these — not create a parallel universe. The goal is to unify standards and decision rights, not centralize every build.

EXISTING GOVERNANCE STRUCTURES

Security & Compliance

Risk tiering, RAI reviews, data classification, identity policies

Cloud / IT Governance

Infrastructure standards, environment strategy, deployment pipelines

Low-Code / Power Platform CoE

Maker enablement, data policies, connector governance, environment strategy

Microsoft 365 Governance

Tenant policies, Copilot configuration, access management

Responsible AI Council

Ethics reviews, high-risk assessments, bias monitoring

↕ Connects to ↕

AGENTIC CENTER OF EXCELLENCE

Unifies agent-specific standards, decision rights, lifecycle discipline, and enablement across all existing structures

Principle: The Agentic CoE adds agent-specific capabilities to your existing governance.
It does not replace what works — it fills the gaps that agents create (ownership, lifecycle, decision rights, monitoring).

Operating Rhythm

The cadence that keeps agents healthy

Example

Weekly
Monitoring & intake

- → Review agent health dashboards
- → Triage new agent requests
- → Address incidents and anomalies
- → Check knowledge freshness
- → Review escalation logs

Monthly
Steering review

- → Leadership scorecard review
- → Outcome KPI tracking
- → Adoption and reliability metrics
- → Risk posture assessment
- → Cross-team coordination

Quarterly
Maturity & value

- → Maturity reassessment against targets
- → Roadmap alignment and planning
- → Business value reporting
- → Portfolio review — what to scale, what to retire
- → Investment prioritization

MONTHLY SCORECARD

Outcome KPI • Adoption rate • Reliability (uptime + accuracy) • Risk posture • Cost-to-serve • Learning loop velocity

Section 4
Your 90-Day Play

From pilots to production — start Monday

From Pilots to Scale in 90 Days

A practical roadmap to operationalize agents in your organization

Days 0–30

WHAT & WHERE

Foundation

- Pick 1–2 adoption patterns that match your current priorities
- Name a specific owner (person, not team) for each initiative
- Run the maturity diagnostic across all five drivers
- Identify your top scale-breaker — the one gap that blocks everything
- Audit your tenant — how many agents exist? Who owns them?

Days 30–60

BUILD & GUARD

Stand Up

- Define minimum governance guardrails (ownership, release gates, monitoring)
- Stand up your CoE rhythm — even if it's 3 people and a weekly standup
- Deliver 1–2 agents to production WITH monitoring from day one
- Create a shared intake process for new agent requests
- Start weekly health checks on all production agents

Days 60–90

RUN & DECIDE

Scale

- Treat agents as production services, not experiments
- Measure outcomes — not just usage and adoption
- Run your first leadership scorecard (monthly cadence)
- Review your maturity progress — has your scale-breaker improved?
- Decide what scales next — which pattern, which team, what investment

You don't need a bigger model.

You need a better operating model.

Resources & Next Steps

Continue your journey with these tools and assets

Agentic AI Maturity Model

Self-assess your organization across five capability drivers.
Use the detailed descriptors in this playbook to place yourself on each level.

aka.ms/AgentMaturityModel

Agent Transformation Stories

Real customer stories. Real agent impact. Discover how organizations across industries are transforming their business with agents and find inspiration for your own AI journey.

aka.ms/CopilotAgentStories

Agentic Transformation resources

Resources to help you on your journey to agentic transformation — building, deploying, and managing agents with Microsoft tools and platforms

aka.ms/agentresources

Microsoft Copilot Acceleration Team

Other resources from the Copilot Acceleration Team (CAT), part of the Copilot Studio engineering team at Microsoft. Our mission is to accelerate the adoption and success of Microsoft 365 Copilot and Copilot Studio.

aka.ms/WeAreCAT

Agents don't scale through technology. They scale through people, ownership, and operating discipline.

You don't need a bigger model.
You need a better operating model.