Architecting an agentic workforce at the EY organization

A strategic framework for enterprise AI implementation

October 2025

The better the question.
The better the answer.
The better the world works.

Shape the future with confidence

Executive summary

As artificial intelligence (AI) transforms the professional services landscape, the EY organization (EY) stands at a pivotal juncture where traditional business models must evolve to incorporate intelligent, autonomous agents capable of augmenting human labor. This article presents a comprehensive framework for architecting an Agentic AI workforce at EY, addressing the technical, governance and strategic considerations necessary for scaling from prototype to production-ready multi-agent systems.

While AI capabilities have advanced dramatically, the fundamental infrastructure of computing, storage and networking remains unchanged. The real transformation lies in how we orchestrate these technologies through sophisticated programing models that can scale from individual agents to complex swarms of thousands, all operating under coordinated governance frameworks. These programing models set the standards for how multiple agents communicate with multiple large language models (LLMs) in a network.

This article is organized as follows:

A. The foundation: understanding agentic systems

B. The EY Agentic Workforce Architecture

C. Agents at work: scaling through complexity levels

D. Governance: the critical success factor

E. Overcoming implementation barriers

F. Practical implementation: tax automation use case

G. The path forward: strategic recommendations

H. Conclusion: embracing the agentic future

A The foundation: understanding agentic systems

An agentic system represents a paradigm shift from traditional software applications to autonomous entities capable of reasoning, decision-making and collaborative problem-solving. At its core, an agent comprises four fundamental components: a model (the natural language understanding engine), data context (the knowledge base), defined inputs (the parameters and triggers) and structured outputs (the deliverables and actions).

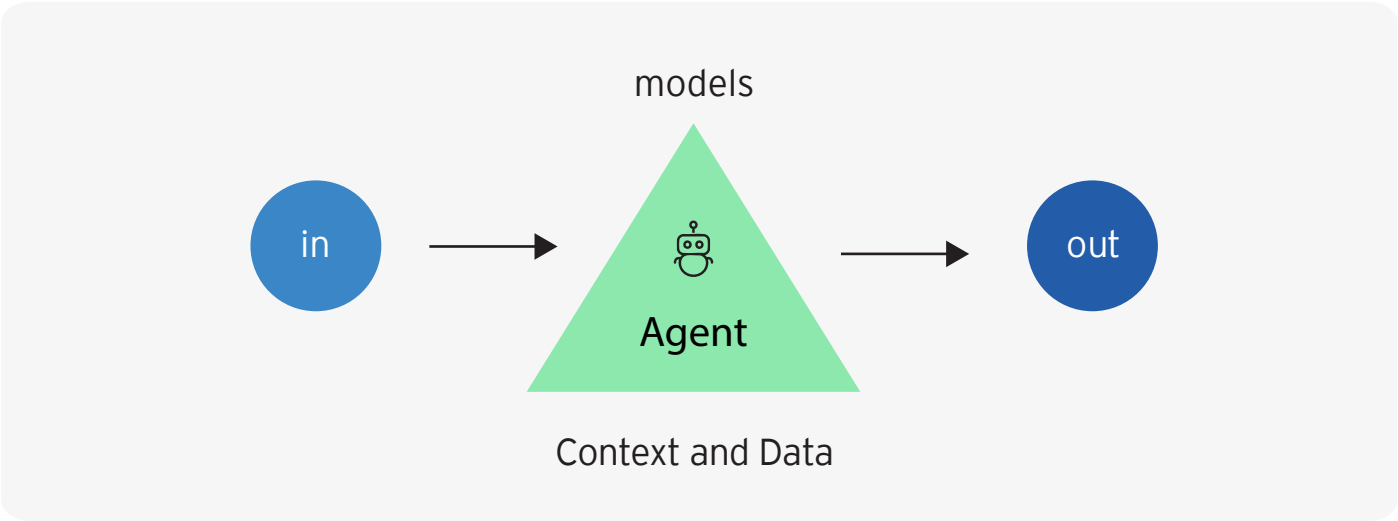

Figure 1: An agentic system

However, the true power of agentic systems emerges when individual agents combine to form multi-agent systems. When two or more agents collaborate toward a common goal, they create a cohesive unit that represents more than the sum of its parts. This basic building block forms the current foundation of EY operational maturity and serves as a billable capability that automates routine, well-defined tasks and complements human expertise.

The challenge is to move beyond individual agents and scale these systems to enterprise levels where hundreds or thousands of agents operate simultaneously under governance frameworks that maintain security, compliance and cost control.

B The EY Agentic Workforce architecture

EY's agentic workforce architecture is built upon five interlocking components that provide both flexibility and control.

1. The agent runtime environment

The EY Agent Runtime Environment serves as the foundational infrastructure where agents operate, communicate and execute tasks. This environment must support both synchronous and asynchronous operations, provide robust error handling and maintain state consistency across distributed operations. The environment acts as the nervous system of the agentic workforce, enabling seamless integration with existing EY systems while providing the scalability needed for future growth.

The Agent Runtime Environment will be flexible enough to not only support EY goals but also be deployable where the client data is located. It will handle the various levels of complexity of running agentic systems so that we can help our clients harness the full power of generative AI (GenAI) within their business process and practices.

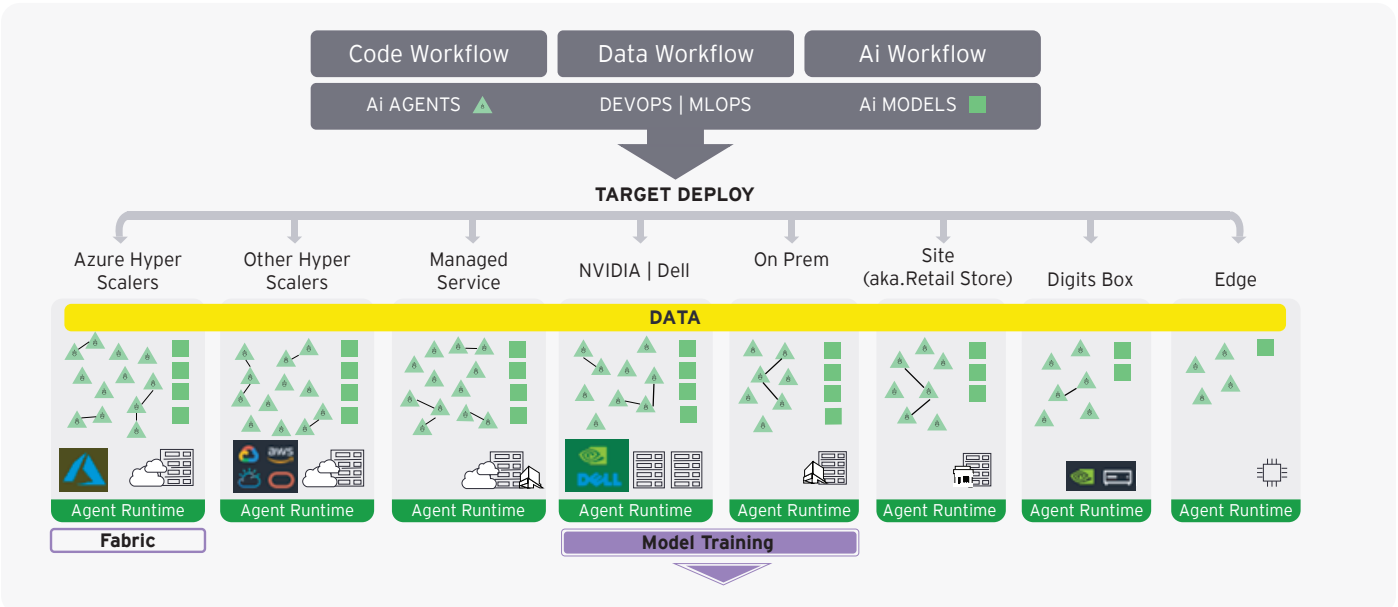

Figure 2: EY Agent Runtime

EY Agent Runtime will be integrated with existing EY systems. For example, Agent Runtime will use data from EY Fabric, which is EY-owned business cloud platform. This platform is centered around enabling EY service lines to bring data close to AI processes like ML Ops, so that this data can be used to train models with specific GPU-accelerated hardware (using partner providers like NVIDIA and Dell as well as our own data centers).

2. Protocol: communication and escalation rules

Agents require sophisticated communication protocols that govern how they interact, share information and escalate complex issues. These protocols define the language of agent interaction, establishing clear pathways for collaboration while preventing conflicts and ensuring efficient resource utilization. The protocols also manage the handoff between agents and human supervisors when situations exceed programed parameters.

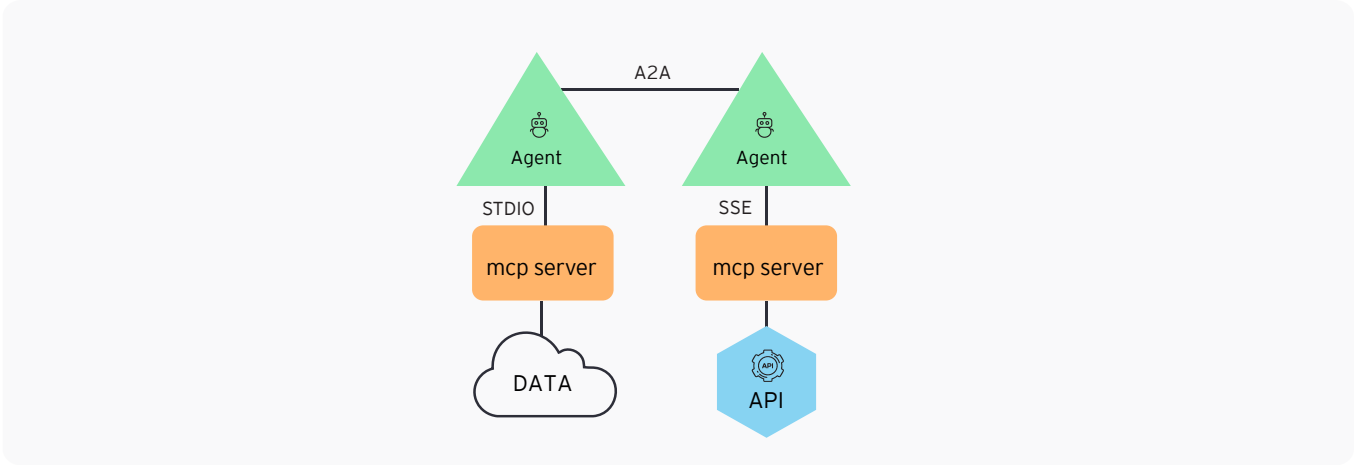

Figure 3: Agent communication protocols

Every technology wave is propelled by the protocols that standardize and connect it. Greenscreen terminals popular in the 1970s and 1980s relied on different terminal protocols such as X11, VT100 and 3270 to connect to mainframes or servers. Similarly, the growth of the web was supported by protocols such as HTTP and DNS. Today's agent ecosystem is following suit with emerging standards like the Model Context Protocol and the Agent-to-Agent Protocol.

These protocols ensure that LLMs receive the appropriate data for reasoning, facilitate smooth task transitions between agents and create a modular foundation for constructing more complex multi-agent systems.

3. Model catalog: domain-specific intelligence

EY model catalog represents a curated library of LLMs and reasoning engines tailored for specific professional services domains. The EY catalog will use domain-language models (DLMs) which are models tailored to specific domains or industries including EY Tax, Assurance and EY-Parthenon, such as healthcare or finance. In some cases, EY teams will use small-language models (SLMs), which are typically smaller in size and require less computational power compared to larger models like LLMs.

In the same way protocols enabling communication between agents, models of all kinds will need to be used by these agents to focus on specific tasks. Therefore, rather than relying on generic AI models, EY catalog will include specialized models for tax preparation, audit procedures, risk assessment and strategic consulting. Each model is fine-tuned for its specific use case and continuously updated based on performance metrics, domain expertise and regulatory changes (for example, changing tax laws). Model alignment to agentic workflow will be a critical activity over the next few years as more use cases are addressed.

4. Patterns: multi-agent workflows

Patterns represent reusable structures that define how multiple agents collaborate to complete complex tasks. These patterns range from simple two-agent collaborations to sophisticated multi-stage workflows involving dozens of specialized agents, including manager agents that direct other agents. By establishing proven patterns, EY teams can rapidly deploy new agentic solutions while maintaining consistency and reliability across different client engagements.

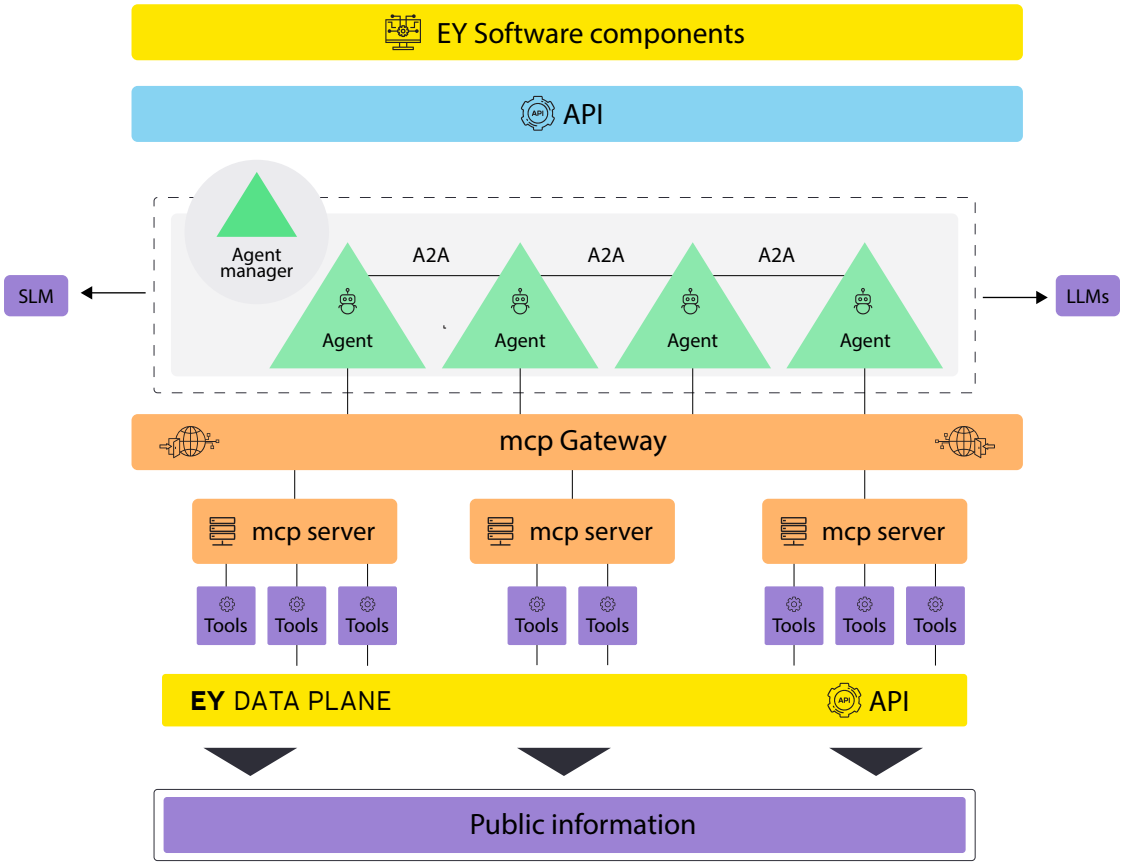

Figure 4: Multi-agent workflows

In all technology that has gone on to support the world's data and computer systems, patterns have formed the approach architects and engineers use to build systems. From the Common User Access guidelines introduced by IBM in the late 1980s to recent, new stack for AI engineering, patterns help shape the direction in which technology is used. We believe that agentic systems will be no different. As new problems are solved, they will be documented into common patterns. These patterns will enable enterprises to scale, and they will also help solve more complex problems.

5. Control rings: governance at scale

The most critical component of the EY model is the control ring architecture. These nested governance levels ensure that agentic autonomy never compromises enterprise compliance or client confidentiality. Outer rings enforce fundamental constraints such as data residency requirements, internet access restrictions and budget limitations. Inner rings manage operational policies like model invocation limits, energy efficiency targets and performance thresholds.

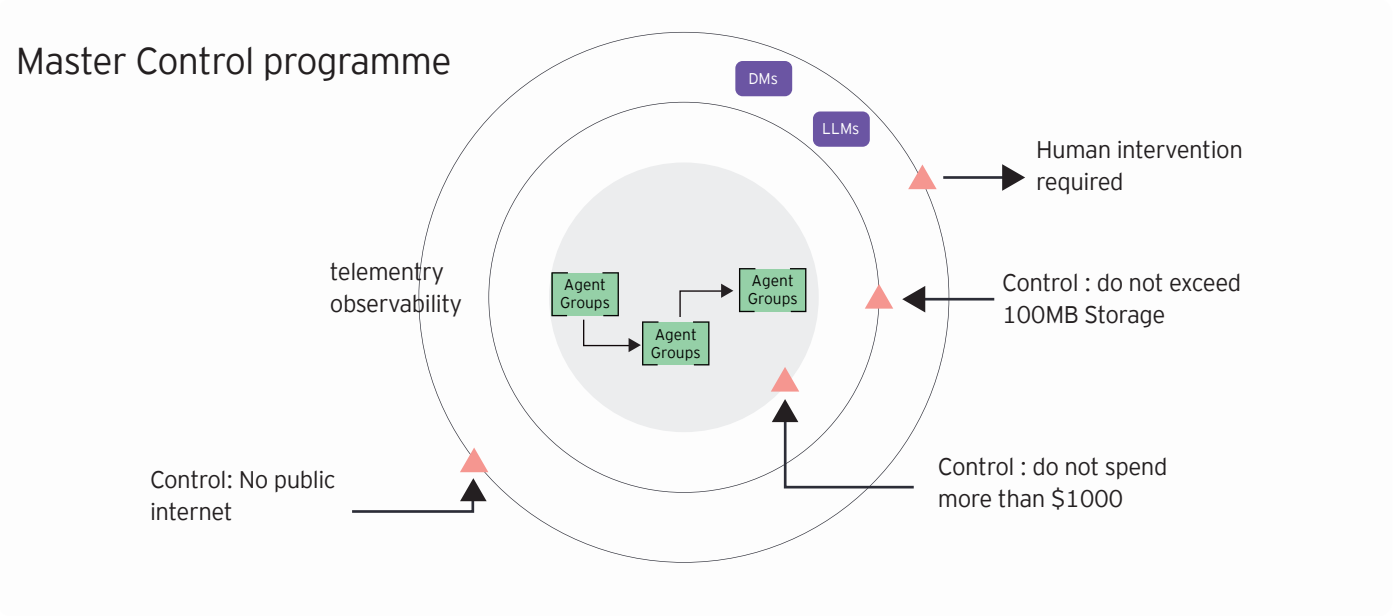

Figure 5: Control rings

A control ring can be wrapped around a single agent or groups of multiple agents. These control rings are discussed further in the section on governance (see section D below).

Agents at work:
scaling through complexity levels

The journey from individual agents to enterprise-scale agentic workforces involves navigating several distinct complexity levels, each presenting unique challenges and opportunities.

Level 1: individual agents

At the foundation level, individual agents handle specific, well-defined tasks. These might include document analysis, compliance checking or data extraction. Each agent operates within clearly defined parameters and produces predictable outputs. The key to success at this level is ensuring robust performance on narrow tasks while maintaining the flexibility to integrate with higher-level systems.

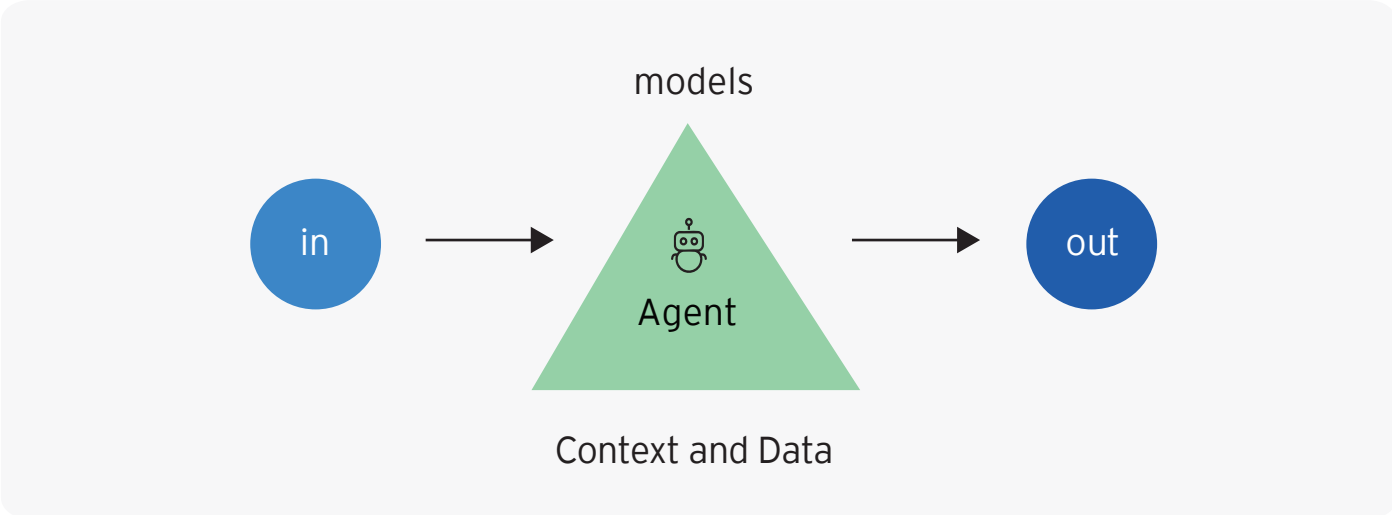

Figure 6: An agentic system

The primary role of an agent is to provide context and data to a LLM, so that the agent can plan and make decisions. Emerging frameworks such as AutoGen, LangChain and Semantic Kernel assist developers in creating functional agents using their preferred programing languages, while also managing much of the complex work involved in the process.

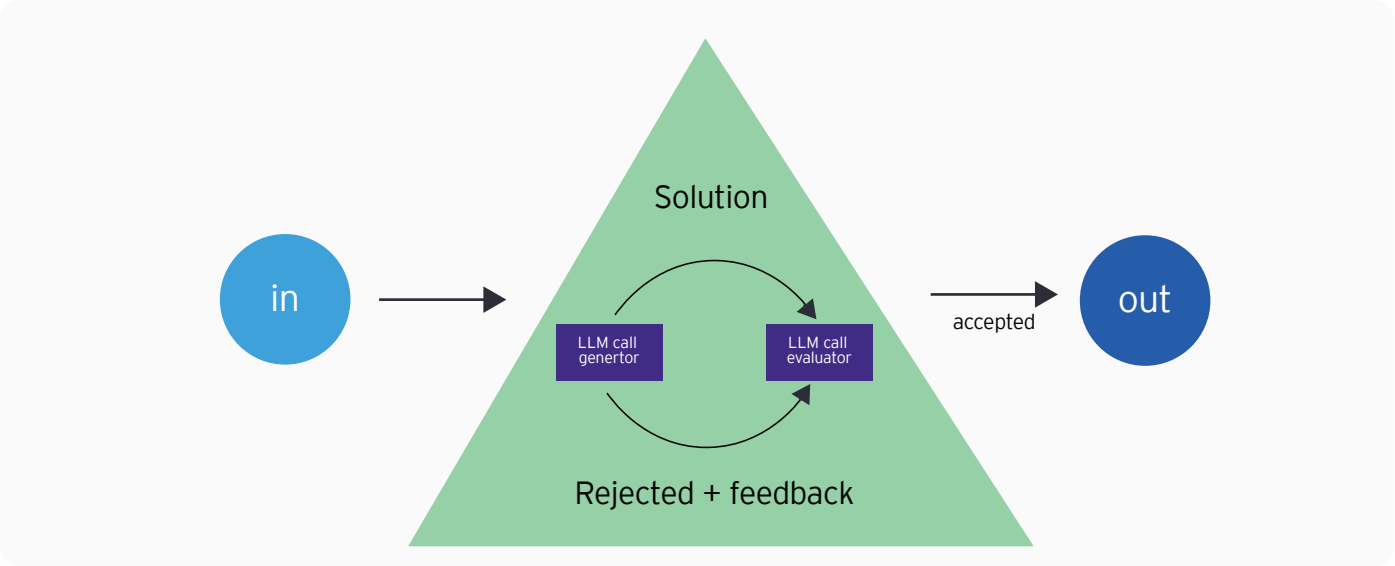

Figure 7: Emerging agentic frameworks

Once each agent can solve a particular task, these agents can be linked together to form a team which can solve more complex problems.

Level 2: agent teams

Teams represent the current operational maturity at EY, typically involving two to five agents working collaboratively on related tasks. For example, a tax preparation team might include agents specialized in data gathering, calculation verification and compliance checking. Teams introduce the complexity of coordination and shared context management while providing significant efficiency gains over individual agent operations.

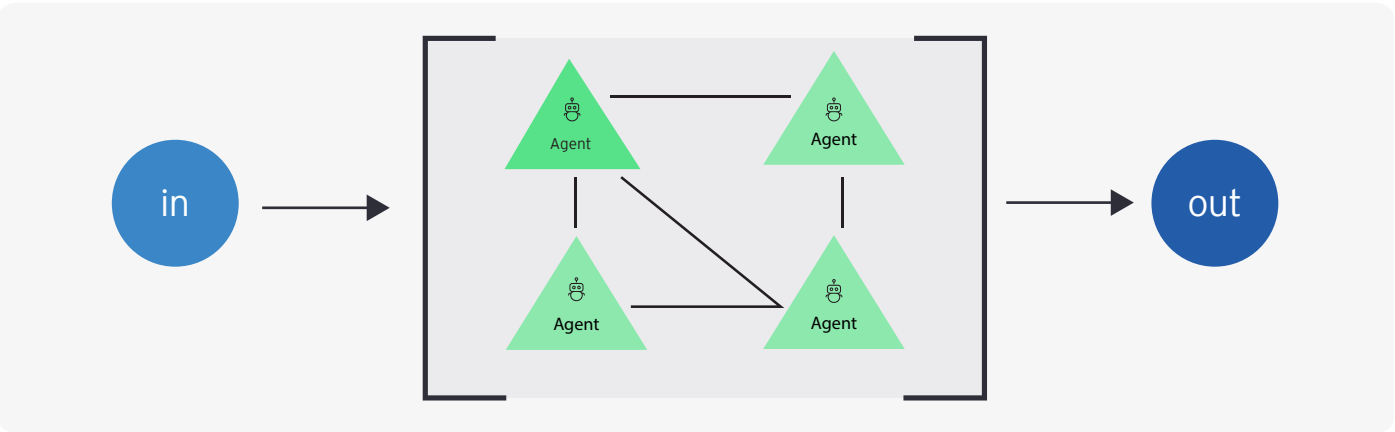

Figure 8: Agent teams

A team of agents can complete a task in a consistent and repetitive way. In a team, there will be an agent acting as the coordinator or orchestrator of the outcome. The team level is the one at which many manual tasks can be replaced with an agentic process, with humans focusing on oversight and exceptions. At this level of agent maturity, people can move up the value chain and help coordinate agentic systems as we are seeing this with vibe coding for developers, where software is written using natural-language prompts to an LLM trained to code, and for self-driving cars for transportation.

We expect that teams will interact with each other leading to the emergence of a pattern. (As mentioned previously, in some cases a pattern may emerge from a single team.)

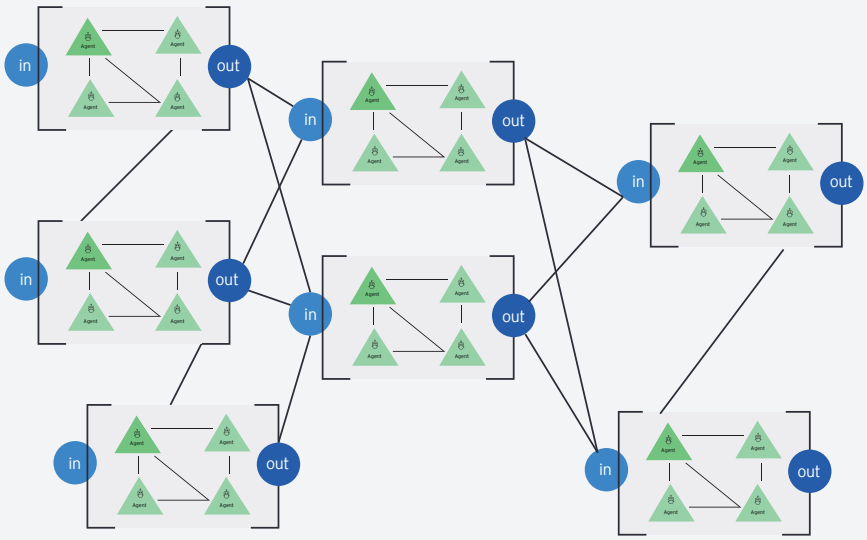

Figure 9: Interaction between agent teams

These fundamental patterns can then interact to form more complex patterns.

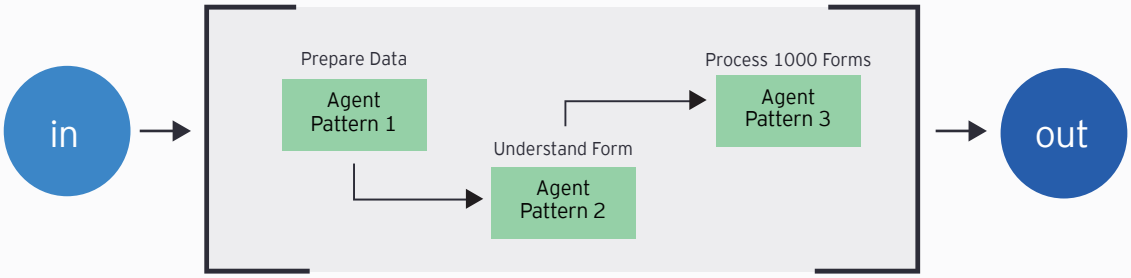

Figure 10: Agent interaction patterns

We expect that complex patterns will lead to the evolution of even more sophisticated, hierarchical patterns that can complete more complex processes and tasks.

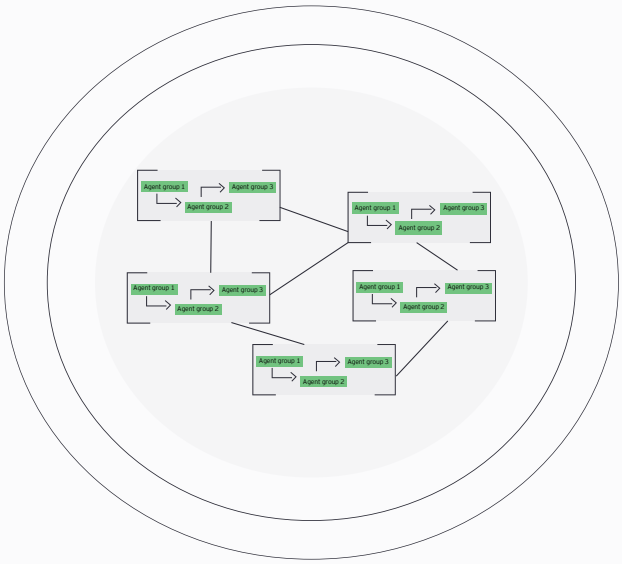

Figure 11: Agent hierarchical patterns

Level 3: organized swarms

Swarms represent a significant leap in complexity, involving many complex patterns with dozens of teams and hundreds of agents operating under coordinated governance. At this level, emergence becomes a critical factor — the system may exhibit behavior requiring continuous monitoring, guardrails and human-in-the-loop controls. Swarms can tackle complex, multi-faceted projects that require diverse expertise and parallel execution across multiple workstreams.

For a swarm, the role of manager or orchestrator is crucial to ensure that all governance requirements are met (see section D below).

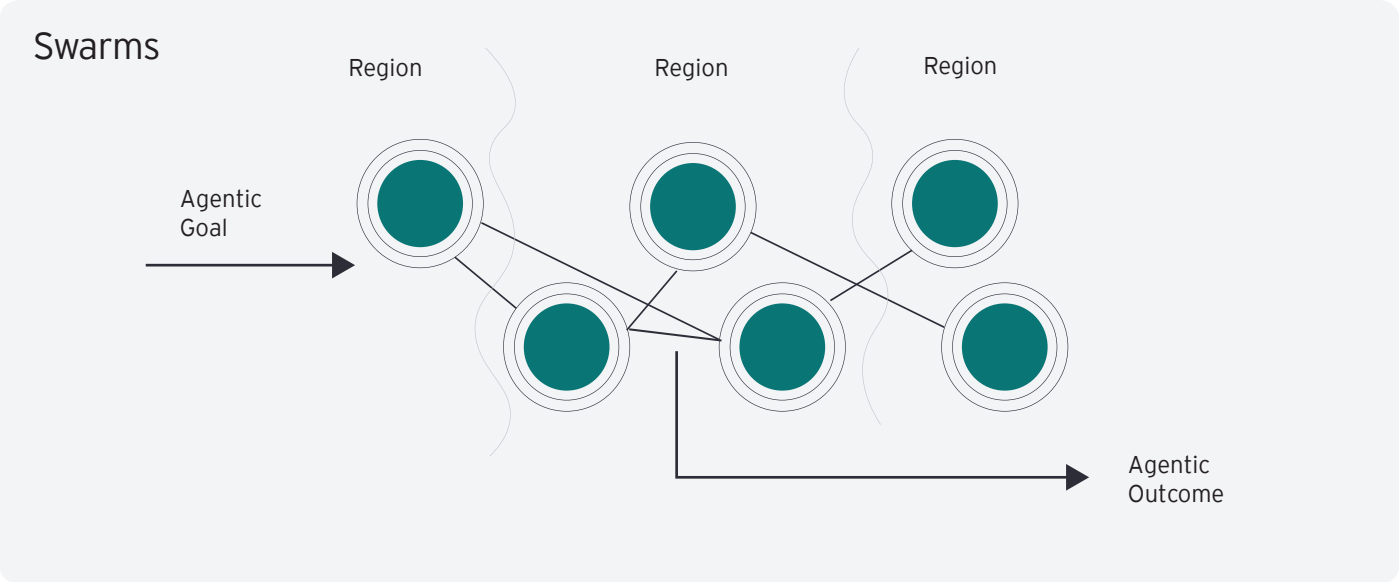

Figure 12: Organized swarms

Level 4: geo-distributed operations

The highest complexity level involves thousands of agents operating across geographical boundaries, managing data sovereignty requirements, language constraints and local regulatory compliance. This level requires sophisticated orchestration capabilities and presents entirely new challenges in areas such as latency management, data synchronization and cross-jurisdictional governance.

Governance: the critical success factor

The sophistication of EY agentic workforce depends on the robustness of its governance framework. Unlike traditional software systems, agentic systems possess a degree of autonomy that requires careful boundary management to prevent unintended consequences. Therefore, the control rings will consist of agents supervised by humans to ensure that all workflows respect security policies, cost limits and enterprise compliance, and client confidentiality.

The purpose of controls is to enable a level of human validation to be applied to a system. The computer industry has been using controls for decades, from the NIST-800-53 controls to the SOC 2 and PCI standards. The AI industry will mature to a point where a set of defined controls will be agreed upon and then implemented. Humans will need to have the tools and interfaces to support human intervention when an agentic process reaches a pre-defined event which is managed by these controls. This area is still maturing, and work continues to be carried out to create an optimal governance process with human intervention. Some essential control rings for this process are discussed below.

Security rings

Security governance forms the outermost control ring, establishing fundamental boundaries around agent operations. This includes network isolation, access controls and data encryption requirements. Security policies might prohibit agents from accessing external networks, require multi-factor authentication for sensitive operations or mandate specific encryption standards for data transmission. Most importantly, given AI's potential for unintended consequences, real-time human review and override protocols will be mandatory for high-stakes decisions.

Cost management rings

Financial governance ensures that agentic operations remain economically viable — running thousands of agents on specialized LLMs could become expensive. Cost rings can enforce budget limits per project, optimize resource allocation across competing priorities or automatically scale down operations when efficiency thresholds are not met by selecting smaller, task-specific models over larger general-purpose ones. This level is particularly critical given the variable costs associated with AI model inference and the potential for runaway resource consumption.

Compliance and ethics rings

Professional services firms operate under strict regulatory requirements that must be maintained even as operations become increasingly automated. Compliance rings ensure that agentic operations adhere to industry standards, client confidentiality requirements and ethical guidelines. This might include automatic auditing of agent decisions, bias detection in AI outputs or enforcement of professional standards in client communications.

Overcoming implementation

The path to widespread agentic workforce adoption faces several significant barriers that must be addressed through thoughtful change management and strategic communication.

Technology fear and resistance

Many professionals fear that agentic systems will replace their roles entirely. EY approach emphasizes augmentation rather than replacement, positioning agents as powerful tools that enhance human capabilities rather than eliminate human involvement. This requires ongoing communication, training and demonstration of successful hybrid implementations.

Client buy-in

For agentic workforces to succeed, clients must trust them. Client buy-in can be achieved by engaging with clients across multiple fronts. First, early pilots should focus on transparency (e.g., explainable AI and audit logs) so that clients understand AI-driven results such as audits and tax filings. Next, clients must know that their data is secure; agents accessing client data must have strict isolation controls (e.g., private LLMs, no external internet access). Furthermore, clients should see that agents demonstrate improvement in speed and quality against agreed, auditable performance metrics. Therefore, it is important to use quantitative metrics to show clients the improvements AI brings. Finally, once clients are comfortable with AI, they should be allowed to tweak agents' behavior (for example choosing conservative vs. optimized tax positions) which could enhance adoption.

Legacy system integration

Existing technology infrastructure at many clients presents integration challenges that can slow agentic adoption. EY framework addresses this through flexible APIs, gradual migration strategies and compatibility levels that allow agentic systems to work alongside existing tools and processes.

Regulatory and compliance concerns

Professional services operate in highly regulated environments where new technologies face scrutiny from regulators and clients alike. EY governance framework provides the transparency and auditability needed to satisfy regulatory requirements while maintaining the flexibility needed for innovation.

Practical implementation:
tax automation case study

To illustrate these concepts in practice, consider the design of an agentic system for tax automation processes:

Single agent structure

A tax preparation agent would include:

- **Model:** A specialized LLM trained in tax regulations and procedures
- **Data context:** Client financial records, previous tax filings and regulatory updates
- **Inputs:** Raw financial data, client information and specific tax year parameters
- **Outputs:** Completed tax forms, compliance reports and exception notifications

Three-agent pattern

A complete tax automation system might involve:

- **Data gathering agent:** Collects and validates client financial information.
- **Calculation agent:** Performs tax calculations and optimization analysis.
- **Review agent:** Conducts compliance checking and quality assurance.

The workflow would involve the data gathering agent preparing clean, validated datasets, the calculation agent processing these inputs to generate tax filings and the review agent ensuring accuracy and compliance before final submission.

Governance rules

Essential governance for this system would include:

- No retention of client data beyond contractually agreed and legally required retention periods.
- Requirement for human review of any filing claiming refunds over a pre-determined amount.
- Compliance with regional data protection regulations.
- Automatic escalation of complex tax situations to human experts.

The path forward: strategic recommendations

Successfully architecting an agentic workforce at EY requires a phased approach that balances innovation with risk management:

Phase 1: foundation building (months 1–6)

Focus on establishing the core framework and developing initial agent patterns for well-defined use cases. This phase should emphasize governance framework development and integration with existing systems.

Phase 2: pattern expansion (months 7–18)

Scale successful patterns across additional service lines while developing more sophisticated multi-agent workflows. This phase should include extensive testing of control ring governance and refinement of pricing models.

Phase 3: swarm deployment (months 19–36)

Deploy complex multi-agent systems capable of handling end-to-end client engagements with human-in-the-loop oversight and exception handling. This phase requires sophisticated monitoring and management capabilities.

Phase 4: geo-scale operations (years 3+)

Implement globally distributed agentic systems capable of handling complex, multi-jurisdictional engagements while maintaining local compliance and optimization.

Conclusion:

embracing the agentic future

The transformation to an agentic workforce represents more than a technological upgrade. It requires fundamental changes in how EY conceives service delivery, client relationships and professional development. Success depends not just on technical implementation, but on cultural adaptation, governance sophistication and strategic vision.

The companies that will lead in the AI-driven professional services landscape are those that master not just individual AI capabilities, but the complex orchestration of multi-agent systems operating under robust governance frameworks. EY five-pillar approach provides a foundation for this transformation, but success depends on execution, continuous learning, and the courage to reimagine traditional service delivery models.

As we stand at the threshold of this transformation, the question is not whether agentic workforces will reshape professional services, but how quickly and effectively organizations can adapt to harness their potential. For EY professionals, the path forward is clear: embrace the complexity, master the governance and lead the industry toward a more intelligent, efficient, and scalable future.

The agentic workforce is not a distant vision — it is an immediate opportunity for organizations bold enough to architect their future today.

EY exists to build a better working world, helping to create long-term value for clients, people and society and build trust in the capital markets.

Enabled by data and technology, diverse EY teams in over 150 countries provide trust through assurance and help clients grow, transform and operate.

Working across assurance, consulting, law, strategy, tax and transactions, EY teams ask better questions to find new answers for the complex issues facing our world today.

EY refers to the global organization, and may refer to one or more, of the member firms of Ernst & Young Global Limited, each of which is a separate legal entity. Ernst & Young Global Limited, a UK company limited by guarantee, does not provide services to clients. Information about how EY collects and uses personal data and a description of the rights individuals have under data protection legislation are available via ey.com/privacy. EY member firms do not practice law where prohibited by local laws. For more information about our organization, please visit ey.com.

ey.com