Department for Science, Innovation & Technology

Government Digital Service

AI Insights

Agentic Workflow

What are workflow systems?

Workflow systems define, automate, and manage sets of tasks and processes within an organisation. These systems provide the tools to design and execute business processes by coordinating tasks, routing information between people and systems, and ensuring that every step is completed in a controlled and predictable manner.

They attempt to streamline business processes by automating the flow of tasks, ensuring that work is completed correctly, on time, and in accordance with defined rules. This leads to increased efficiency, transparency, and the ability to quickly adapt to changes in business requirements.

Legacy systems

Legacy workflow systems are designed to manage business processes using fixed, sequential steps. They typically rely on manual intervention, with tasks and approvals often handled through static forms or direct human input. These systems have limited automation and integration capabilities, making them less adaptable and harder to scale compared to modern, dynamic workflow solutions.

Workflow frameworks reflect this fixed, sequential model. They often require specialised knowledge so bridging between steps and passing information between steps can quickly become a cumbersome process.

Closely related to workflow systems, business process modelling (BPM) applications have a similar tendency to become complex quite quickly, and are often proprietary - only available within a single vendor ecosystem. These are difficult problems to solve.

What is agentic workflow?

Agentic workflow is a novel approach in which autonomous artificial intelligence (AI) agents manage, coordinate, and execute tasks within business processes. In an agentic workflow, each component or agent operates independently, making real-time decisions based on predefined rules, data inputs, and contextual insights. This allows the system to adapt dynamically to changes or unexpected events without relying on human intervention.

Similarly, microservices architectures illustrate distributed coordination though modular, independent components which interact through defined interfaces. However, these architectures remain constrained by static rules and predefined interactions.

Agentic workflows can enhance these models by embedding autonomous decision-making directly into each agent, enabling dynamic collaboration, adaptive task management, and spontaneous responses to unexpected events.

Agentic workflows commonly employ multiple agents to complete complex tasks, which the agents are able to achieve through collaborative information sharing. The interactivity of multi-agent systems tends to promote visibility of processes. That is, we can map the workflow pathways and optimise them, and also see any bottlenecks or congestion in the system and raise alerts appropriately. This also enables the system to scale and adapt as business requirements evolve.

In the context of agentic AI, instead of just following step-by-step, hard-coded instructions, the AI takes on a more active, self-directed role within pre-defined parameters.

Autonomy

Agentic AI systems are designed to make decisions on their own. An agentic workflow encapsulates the steps the system follows. From understanding a high-level goal to breaking it down into sub-tasks, choosing the best actions and then executing them, all without needing explicit instructions at every step.

Dynamic task decomposition

In an agentic workflow, the AI isn't limited to a fixed, linear process. It can dynamically decompose complex tasks into smaller, manageable actions. For example, if tasked with planning a holiday, the agent might autonomously decide to research destinations, check travel prices, schedule accommodations, and then adjust its plan as new information becomes available.

Agentic workflow lifecycle

A typical agentic workflow involves:

Planning

The AI interprets the overall goal and devises a strategy or sequence of actions.

Execution

It then carries out these actions, which might involve interacting with external systems such as databases, application programming interfaces (APIs), or other agents.

Iterative and adaptive processes

Unlike rigid workflows, an agentic workflow is iterative. It incorporates feedback loops where the AI re-evaluates its actions and outcomes, refines its strategy, and repeats the process until the goal is achieved. This makes the system resilient to unexpected changes and capable of handling complex, evolving tasks.

Monitoring and Feedback

The system continuously monitors its progress and the results of its actions. If something isn't working as expected or conditions change, it can adapt its plan on the fly.

Practical Examples

Automated business processes

This is an agentic AI system that manages supply chain logistics. Instead of following a fixed set of instructions, it may detect delays, re-route deliveries, adjust inventory orders, prioritise perishable items, offer green discounts for deliveries in your area on

common delivery dates, and communicate with suppliers autonomously, all as part of its agentic workflow.

Personal assistants

An advanced personal assistant might use an agentic workflow to plan your day, scheduling meetings, making reservations, sending reminders, and rescheduling activities if conflicts arise, all by assessing real-time data and knowing the history of your preferences. This becomes even more powerful when agents collaborate to schedule meetings across multiple busy calendars.

Autonomous travel planners

This agent, or multiple agents, can autonomously schedule flights, hotels, and other arrangements for a person or team. This is non-trivial, and includes goal interpretation, dynamic task decomposition, multiple points of synchronisation, and communication with external systems or agents. This is a complex behaviour with a very high degree of utility.

The autonomy and flexibility of agentic workflows are made possible by integrating large language models (LLMs) and their powerful comprehension and planning capabilities. However, we must provide the capabilities to these systems. In the above logistics example, we provide an agent controller with various functions it may call and tools it may use, along with natural language descriptions of these capabilities.

We will also typically provide the system with a memory capability so that it can record how far along it is in the process and be better able to decide what to do next. To do this, the system will consider its current state (using data), where it needs to be, what capabilities it has at its disposal, including its ability to plan and route, and finally determine and execute the next best course of action..This is a very powerful and flexible paradigm, combining multiple intelligent processes alongside human architecture and design, and the coding of capabilities. When enough of these modules are combined we may find ourselves with some very powerful workflow systems.

Risks

There is, of course, no free lunch. Any novel technology will bring new risks and concerns to the table. Here are a few examples of areas we must consider when designing and developing agentic workflow solutions:

Transparency

Simple linear systems are relatively easy to follow but complex traditional systems can become rather convoluted as they grow and the number of possible branches or routes they follow increases. However, generally speaking, we know the steps taken under all conditions and can follow the path from start to finish.

Agentic systems may be less transparent. Given their autonomy, their independent decision-making, and their collaboration with other autonomous agents, it might be more challenging to follow the path of execution. This is a well-known issue with autonomous systems.

Fortunately, there are many different solutions to this, such as deliberately maintaining an audit trail which can be used to profile and optimise our interactions.

Unintended behaviours

With linear systems there is a hard-coded, deterministic path. Agentic systems, again due to their autonomy, are simply not guaranteed to make the best decision at all times. This may show itself especially in rare or complex scenarios where the agentic system has less information or its pathway selection is obscure.

We have means to mitigate this, mostly by being aware of it as a design and development responsibility. Agentic workflow must be designed, developed, configured and tested appropriately in order to achieve system objectives.

Given the possibility of unintended behaviours, it is the responsibility of the system designer to ensure that the profile of the system under expected conditions is known and controlled. Testing specifically for these cases should be mandatory. We must also understand how the system behaves outside these expected conditions and build effective guardrails. Regular reviews of system performance should be scheduled to mitigate this issue.

Underlying models

The effectiveness of an agentic workflow often hinges on the performance of the LLM in use as this governs the agent's decision-making. These models can have biases, produce hallucinations, or simply make errors, which, when compounded across multiple agents, can undermine the system's reliability.

Data validation, model profiling, and testing are essential to know as they manage the limits of these systems. We should also note that if the agent's LLM is changed, extensive testing must be undertaken to check what implications that has for the system.

Refer to our agentic AI Insight, and agentic RAG Insight for more information.

Is agentic workflow the future?

Agentic workflow is a new, different type of workflow,and not likely to replace traditional workflow systems for many reasons. Not least of these is cost; at this time, it simply does not make financial sense to replace all legacy workflow systems, especially those that are working perfectly.

The significant impact of agentic workflows on energy consumption must also be considered, as the frequent use of LLMs are a key driver of increased demand. Agentic workflows rely on LLMs for contextual interpretations and decision making, requiring repeated, high-intensity computations. Unlike traditional computing systems which handle predictable tasks with consistent resource allocation, agentic workflows consistently adapt by dynamically distributing tasks and frequently using resource-intensive LLMs.

Although the computation intensity of LLMs inherently increase energy use, the intelligent task delegation within agentic workflows can offset this by efficiently assessing workloads, minimising redundant computations, and reducing idle resource issues. This allows them to balance their high energy demands with optimised resource management.

As for agentic RAG, this is another intelligent instrument in our toolset. When the brittleness or complexity of our traditional systems becomes a burden, this may be best alleviated by an agentic workflow.

Conclusion

In summary, agentic workflows represent a shift from rigid, predetermined processes to autonomous, intelligent systems. By enabling AI agents to interpret high-level goals, devise action plans, execute tasks in real-time, and continuously refine their approach based on feedback, these workflows offer tremendous utility.

Looking ahead, traditional linear workflows will continue to offer value for a great many use cases. We will likely have traditional and agentic workflows operating side-by-side.

We will also likely see an increasingly seamless integration between them, with conventional systems combined with the intelligence of agentic workflows.