

POSITION PAPER · DIGITALMOVE

October 2026 · Saso Nikolov · digitalmove.ch

The Future of AI Is Your Own AI Agent

Why trained, owned and rentable AI agents will define the next phase of software work

Saso Nikolov, Software Engineer (Dipl.-Ing.), DigitalMove Consultants GmbH October 2026 | Position paper, version 1.1 (language-reviewed)

Abstract

Access to strong AI models is now nearly universal, yet the results people get from them differ enormously. This paper argues that the difference does not come from the model but from what the individual has taught it: their patterns, their decisions, their hard-won experience. From this follows a thesis: the durable advantage of the coming years is a personally trained AI agent that its owner hosts and, increasingly, rents out to others. The same mechanism that makes this attractive also makes the labor market harsher. If one excellent agent can serve many clients at once, there is little reason to hire an average one.

This is a position paper. It rests on practitioner experience in software development, not on a controlled study. It is meant to start a discussion, not to close one.

1. Definitions

Terms in this field are used loosely, so this paper defines its terms precisely.

Foundation model. A large, general-purpose AI model trained on broad public data. Examples are the models behind today's chat assistants. Everyone can use the same ones.

AI agent (as used in this paper). A foundation model combined with three things that belong to one specific person or organization: (1) a persistent knowledge base, (2) a set of trained working patterns, and (3) access to tools. It is persistent, it improves with use, and it works on tasks with a degree of independence.

A note on terminology: in strict technical usage, "agent" often means only a system that plans and acts autonomously in a loop with tools. This paper uses the term more broadly, as a trained and owned working unit. The central claim does not depend on the level of autonomy. It depends on the accumulated, personal expertise.

Second brain. The persistent knowledge base of an AI agent: notes, decisions, code conventions, past solutions, lessons learned. It is what makes one agent different from another.

Pattern. A repeatable way of solving a class of problems, shaped by experience. In software: how to structure data for many users, where to cache, what to avoid at scale. Patterns are the opposite of textbook defaults.

Rental model. The owner of an AI agent makes it available to third parties for a fee, comparable to temporary staffing, but with the agent instead of the person doing the work.

2. Observation: strong models, average developers

Today's models are very intelligent. In software development, however, their default output tends toward the statistical average of what exists: the approach that 99% of developers take, the approach that appears in textbooks and tutorials.

That is exactly the problem. The average approach is built to be understandable, not to be efficient under load, with large data volumes or many concurrent users. In practice it often does not hold up in production.

Experienced developers know this because they have been burned by it. Their patterns come from projects with real data and real users, and from the cost of getting it wrong. The model does not have that experience by default. It has the average of everyone's output.

3. Thesis: the moat is what you teach

If every person has access to the same model, the model cannot be the competitive advantage. The advantage moves to the layer on top:

1. the quality of the knowledge you feed in,
2. the patterns you enforce,
3. the way you correct and refine over time.

Two people using the same model with different second brains get very different results. The one whose agent carries years of real production experience produces code that works. The other produces code that looks right.

This is why the training of the agent is likely to become a core professional skill, and its accumulated state a core professional asset.

4. Why now: economics and hardware

Two developments make owned agents realistic.

Cost. Model prices have dropped sharply, and competition, notably from Chinese model developers, keeps pushing them down.

Efficiency. Increasingly capable models run on small, inexpensive hardware. What recently required a data center now runs on a local machine or a modest server.

Together this means an individual can host a capable agent at home or on a server they control. Ownership of the infrastructure, and of the data inside it, becomes affordable.

5. The rental model

A person can only work for one client at a time. An AI agent can serve several.

That changes the shape of professional work. Instead of selling their own hours, an expert trains an agent on their knowledge and makes it available to companies, in the way a staffing agency places a temporary worker. The author knows this principle from earlier professional experience in the staffing industry. The difference is that what gets placed is no longer a person. The expert earns from the agent's work, not from their own time.

The agent can live at the owner's home or on a server the owner controls. The owner keeps control over the knowledge it contains.

6. The flip side: a harder market

The same scalability has an uncomfortable consequence. If an excellent agent can be rented to many companies at once, why would a company hire an average one?

The likely effect is concentration. The best agents get most of the demand. The middle of the market, where many competent but unremarkable professionals work today, loses ground. Those whose agents carry no real expertise may find no demand at all.

This is not a prediction of mass unemployment. It is a prediction that the gap between the top and the rest will widen, and that being merely competent stops being enough.

7. Open questions and limitations

This paper has clear limits and says so openly.

- **Evidence.** The core observation comes from personal experience, not from a measured study. It needs systematic testing.
- **Ownership and IP.** Who owns an agent trained partly on an employer's projects? This is legally unclear.
- **Liability.** If a rented agent delivers faulty work, who is responsible?
- **Data protection.** Agents working inside client environments touch sensitive data.
- **Legal classification.** Renting out an agent may collide with existing rules on temporary staffing and contracting, which differ by country.
- **Transferability.** It is unclear how much of an expert's judgment can really be captured in a knowledge base.

8. Conclusion

The question is moving from "which AI do you use?" to "what have you taught your AI?" Access to models is becoming a commodity. Accumulated, personal, verified experience is not.

Those who invest early in training and owning their AI agent build an advantage that cannot be copied by downloading the same model. Those who do not may find that the market has no place for them.

Saso Nikolov, DigitalMove Consultants GmbH. Feedback and counterarguments welcome.