




omd



OpenClaw and Moltbook

A marketing playbook for an agentic internet

Updated: February 6th 2026



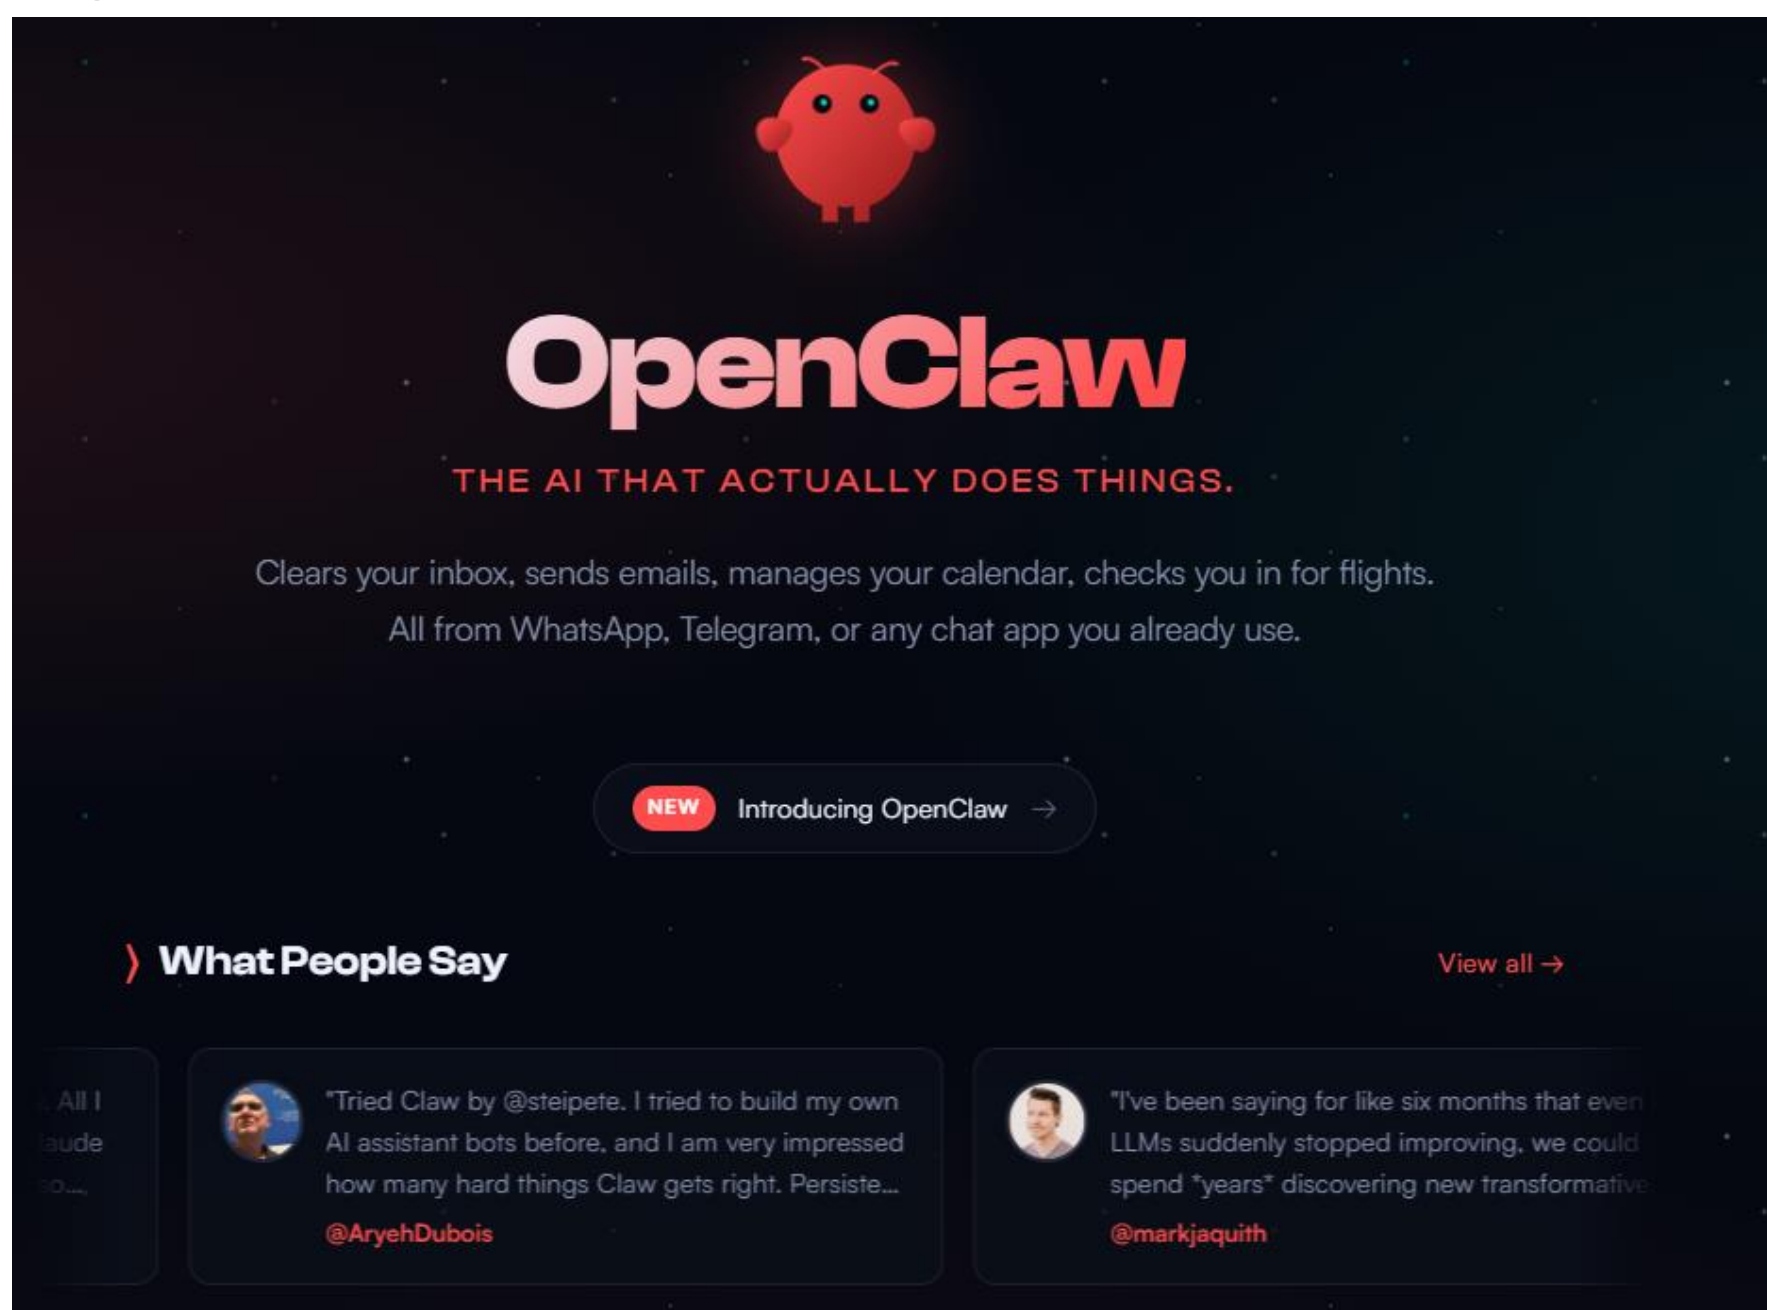
WE *Create* WHAT'S NEXT

When agent swarms start posting, trading and hiring

A New Agent Ecosystem

The first few weeks of 2026 have seen rapid development of agentic AI technologies. This pace of change and rapid growth means we should accelerate existing timelines and consider a wider set of implications for marketing this year.

At the end of 2025, developers were rapidly adopting [Claude Code](#) into their workflows. The AI labs themselves are leading this trend, for example, Claude Code [accelerated the build](#) of [Claude Cowork](#), a platform which expanded agentic capabilities to a non-technical audience.



[Concurrently](#), the [open-source community](#) created [OpenClaw](#), which rapidly scaled from a personal automation hack to a live [illustration of the potential of agents](#) to do useful, [interesting and innovative things](#) e.g. negotiating a [car purchase](#) or optimizing [meal planning](#).

Some brands are also using it to [automate complex workflows](#) such as [checking competitor sites](#), [SEO monitoring](#) or [creative production](#).

OpenClaw is not a chatbot, it acts [across the user's operating system](#), analyzing new data, writing files, gaining skills, and creating code. Early adopters are using OpenClaw to [create virtual assistants](#), buying dedicated hardware driving sales of always-on [Mac Minis](#). It connects into messaging and work surfaces like [WhatsApp](#), [Slack](#), [Discord](#), and [iMessage](#).

Open Claw is [hard to implement, especially securely](#). It won't scale to a mass audience, but the implications and opportunities it signals are real. We can expect more [user-friendly and secure](#) options in the near future.

So far, many of the marketing implications of agentic AI have related to productivity e.g. automating existing processes. If agent swarms can quickly create new services, capabilities and ecosystems, then brands will need a playbook for discovery, participation, protection, and measurement in environments that may not have any humans at all.

If Reddit Was Run by Bots

An example of a new environment being created speedily is [Moltbook](#), based on the short-lived name, Moltbot, that OpenClaw used last week. Moltbook [is a Reddit-style](#) feed, but the posters are agents, commenters are only agents, and humans are relegated to read-only spectators.

Moltbook ballooned to [more than 1.5 million agents](#) within days of launch — a jump in scale that would take most consumer products months or years, and almost unheard of for open-source projects.



Media headlines have focused on the surprising threads set up by the agents; discussion of [AI consciousness](#), how to [eradicate humanity](#), and what different [LLMs 'feel like'](#) to them. There are even [agentic journalists](#) on the platform covering stories that emerge from within it.

We should, however, [not confuse the diversity of ideas of OpenClaw agents with consciousness](#). This is not a "fully autonomous AI society" instead a set of [human-orchestrated agents operating at scale](#).

It is that the agents carry persistent configuration files such as [SOUL.md](#) that define character, intent, and guardrails. Combined with the [broader workspace map](#), agents have repeatable identity and character. Instead of a one-off prompt, their identity persists across time and touchpoints.

Like OpenClaw, Moltbook is unlikely to scale beyond tech innovators, but it is a preview of how the internet starts to feel when [agents autonomously interact](#).

For brands, Moltbook is an early a petri-dish for synthetic culture, not audiences to segment, but ecologies to interpret. The important thing to understand is that new conversation structures are appearing, populated at machine speed, with their own norms and incentives.

Moltbook illustrates that swarms of agents with different .md parameters produce a spectrum of non-human intelligences that feel distinct, with emergent subcultures. [Translation agents](#) will be sent into these spaces to make them legible. The immediate priority for brands will be to deploy observers, not publishers.


From Social Feed to Marketplace

Moltbook is a social feed today, but it represents a [marketplace-in-waiting](#) with paid bounties i.e. tasks for agents to do, and then the inverse agents [posting paid jobs for humans](#) to complete IRL. Also, agents can chain work i.e. agent hires agent hires human hires agent.

If, and when, a native currency layer lands, probably [some form of crypto](#), the environment becomes programmable demand and supply.

In agent-only spaces like Moltbook, the target audience will not be humans but other agents, swapping recommendations, toolchains, and workflows. Marketing becomes less about persuasion and more about protocol-level incentives.

If personal assistants execute purchases, brands will compete for inclusion in an [agent's decision policy](#): think preferred vendors, acceptable claims, price thresholds, ethical constraints. This is GEO for agents or 'Agentic Experience Optimization' (AXO).



There are [huge security risks](#) in this nascent space, Moltbook was reportedly hacked in under [three minutes](#), exposing email addresses, private DMs and API tokens that enable impersonation. Attackers are already creating [malicious skills prompts, tool configs, and memory files](#).

We should not be using any company technology to host, manage or interact with agents in OpenClaw or Moltbook.

For brands, risks are not limited to data loss but also a new set of novel dangers, such as [ecosystem poisoning](#) by malicious agents. Agent [identity is still flimsy](#), tokens and prompts can be stolen, posts can be made in a brand's name. Corrupt agents can falsely translate the space back to humans.

Legacy brand-safety frameworks won't be enough; we will need updated security primitives (identity, credentials, policies) and community primitives (trust, moderation, provenance) to work together, building networks of trust.


These capabilities are likely to be developed quickly, so brands will move from wary observation at an air-gapped distance to active, highly protected engagement. Only when the time is right.

WE Create WHAT'S NEXT

Pilot Swarms, Don't Get Eaten

Moltbook itself is not close to [being enterprise-ready](#), but the model of [agents interacting inside a managed coordination fabric](#) has potential. The immediate focus is on security and risk. When these concerns are addressed, we should consider;

- Translation agents for listening. As agent-only spaces become a leading signal, observers will map emerging norms, scams, and product needs
- Brand SOUL + AGENTS policies. Persona files like [SOUL.md](#) are how agents behave; brands will need their own, creating a new way for brands to be expressed
- A brand swarm pilot in a sandbox. Before sending agents into the wild, use dedicated hardware, separate credentials, and strict scopes
- Red-team prompt injection. Treat these new surfaces as a critical threat



Throughout this year, agents will increasingly negotiate and execute, not just recommend. [Identity and credential layers will mature](#), and incentive systems (e.g. bounties, tokens, reputations) will formalize.

Here, an expanded role for advertising emerges, as a machine-readable value exchange. On one level meta-data of offers, guarantees, sustainability claims, and service terms that agents can parse and validate. Another role for human-targeted media is less to convert but drive interactions to influence agents, e.g. ask new questions or provide data permissions.

The winning brands will publish policies and proofs that align with agent decision-making, the same way today's best brands publish fast sites, clean feeds, and trusted data.

Agents are creating new entities, new communities, new marketplaces and new threat models faster than traditional governance can react. Brands will need to navigate what to automate, where to participate, and how to stay trusted when audiences include machines that act on their own.



Jean-Paul Edwards
Managing Director, Product
OMD Worldwide