How to put an AI agent on your old laptop with maximum security

All you need is five distinct layers of isolation, each designed to prevent a specific and unique threat

Pavel Gurov · ORCID 0009-0001-0152-8056
Written 17 March 2026 · 1141 words · gurovdigital.com/research/ai-agent-security-five-layers

Agent security · AI security · AI infrastructure · Autonomous agents

How to install an AI agent on your old computer in a COMPLETELY secure way.

Corporations in the United States are currently paying six thousand dollars for a secure setup of OpenClaw, the new-generation AI. Below I will explain how to build enterprise-grade security for an AI agent. This is entirely unique information that has not been published anywhere.

Most people were frightened off by two things

As you know, over the past month the technical world has been obsessed with the OpenClaw revolution. This is an AI that actually does things, an AI agent that is the next revolutionary step after all the GPT chats. And of course many people want to get into the new future, but most were afraid to install OpenClaw.

First, most people thought that you need to buy a six-hundred-dollar Mac Mini. That is not the case. I installed OpenClaw on my old 2016 MacBook. As you can see, it stands propped up like a tent so that it does not overheat, because it is so old that there is not even a cooler inside it, no fan. And I stand it like a tent because it does not even need a screen; I control it remotely from my main Mac through the Terminal program. Yes, OpenClaw needs an isolated machine - that is precisely for security, you cannot put it on your main machine, because otherwise attackers can hijack it.

The second reason concerns security. Most people were frightened that OpenClaw could be very dangerous, and there were indeed several stories - for example, an Israeli cybersecurity company found a terrible vulnerability. But as I wrote in my previous post, this is a vulnerability of absolutely any agent, including the code assistant beloved by all white-collar workers[1].

So the good news is that installing a new-generation agent has now become simpler.

I wiped the machine before I let the agent anywhere near it

The most important thing: listen to me, I am now going to tell you my personal experience. I installed OpenClaw on my old MacBook. Yes, I had to spend some effort, because I made a special account. First, this MacBook was reset to factory settings, so that none of my personal pho-

tographs or data or Chrome sessions were left on it. With prompt injections it is possible to steal your session. That is not even login and password - it is worse. It means that you will not even notice anything, and you can read about session hijacking in more detail elsewhere.

So, I reset it to factory settings and removed the Apple ID entirely, meaning it is not needed there at all. In fact there is no Apple ID on it, because all applications will be installed through the terminal. Next, what I did: I created a special admin account for the agent with minimal rights, without root access, so that it cannot destroy the system.

The old MacBook, wiped to factory settings, that the agent runs on.

After that I installed not OpenClaw but Nemo-Claw. What is that? Yesterday Peter Steinberger, the creator of OpenClaw (a solo founder, like me), joined forces with NVIDIA and with its head, and

they presented yesterday[2].

> *Mac and Windows are the operating systems for the personal computer. OpenClaw is the operating system for personal AI.*
>
> 
> **Jensen Huang**
> Founder and chief executive, NVIDIA
>
> Photograph by Peter Dasilva, cropped, CC BY 4.0

They announced that the NVIDIA corporation has built an additional protective shell specifically for corporations. This is still, of course, not one hundred per cent protection. One hundred per cent protection does not exist anywhere in the world. But it is a significant step forward, because it creates a shell and mitigates the threat of malicious code execution[3].

> *With NVIDIA and the broader ecosystem, we're building the claws and guardrails that let anyone create powerful, secure AI assistants.*
>
> 
> **Peter Steinberger**
> Creator of OpenClaw
>
> Photograph by dotconferences, cropped, CC BY 3.0

After that you attach Slack to it, where instead of colleagues your various agents live (not only OpenClaw, but also Manus and Claude Cowork - but I will write about that in future posts).

Each layer stops a different thing

Layer 1: Identity isolation. Implementation: a completely wiped Mac with no Apple ID attached. What it protects against: it rules out the theft of authentication tokens. If the agent is compromised and gains access to the macOS file system, it will not be able to extract iCloud session keys, reach your personal photographs, your Keychain passwords or your iMessage history.

Layer 2: Network and hardware isolation. Implementation: a headless server, closed lid, inbound ports closed (using WebSockets Secure for Slack). What it protects against: it rules out external network attacks. An attacker cannot scan the ports on your router and find a vulnerable agent interface, because the server only initiates outbound connections and does not accept inbound ones.

Layer 3: Access control at the operating system level. Implementation: an `aiagentworkspace` account without superuser (root) or sudo privileges. What it protects against: it implements the Principle of Least Privilege. If an attacker breaks out of the Docker container into the host system (macOS), they end up in the sandbox of a restricted user. They will not be able to install a rootkit, alter the firewall configuration or read encrypted system files.

Layer 4: Base containerisation. Implementation: running OpenClaw inside Docker Desktop with hard resource consumption limits. What it protects against: it isolates the file system and the

processes[4]. The agent "sees" only those files that we have explicitly allowed it to see inside the virtual environment, and it has no direct access to the `/Users/pavelgurov/` directory.

Layer 5: Runtime sandbox - NemoClaw's area of responsibility. Implementation: integration of NVIDIA OpenShell (isolation at the level of system calls). What it protects against: it complements Docker. Docker is a static box. If the agent needs to execute a Python script to analyse data, Docker will allow it. NemoClaw analyses the script's intentions in real time. If the script attempts to execute a curl command to covertly send data to an external server, or attempts to read the system file `.env`, OpenShell will block that specific system call without stopping the work of the entire container.

The architecture has no single point of failure

Critical consideration: your architecture has no single point of failure[5]. If an attacker finds a zero-day in NemoClaw's protection, Docker will stop them. If they break through Docker (container escape), the restricted rights of the macOS account will stop them.

Performance risk: on x86_64 architecture processors (your Intel Mac) Docker Desktop already runs inside a lightweight virtual machine managed by the macOS hypervisor. Adding the OpenShell module on top of this stack will increase processing latency. Expect the execution of complex commands through Slack to take 1.5 to 2 seconds longer than usual because of deep packet inspection.

Insight: you have built infrastructure that meets enterprise-grade security standards. This is excessive for testing basic chatbots, but absolutely necessary for autonomous systems that have the right to execute code and access the internet.

REFERENCES

1OWASP GenAI Security Project, *OWASP Top 10 for LLM Applications* - https://genai.owasp.org/llm-top-10/↩

2NVIDIA, *NVIDIA Announces NemoClaw for the OpenClaw Community*, press release, 16 March 2026 - https://nvidianews.nvidia.com/news/nvidia-announces-nemoclaw↩

3Cloud Security Alliance, *NemoClaw Security Assessment: Enterprise Agent Runtime Hardening* - https://labs.cloudsecurityalliance.org/agentic/↩

4Palo Alto Networks, *What Is Container Security?* - https://www.paloaltonetworks.com/cyberpedia/what-is-container-security↩

5NIST, *Defense in Depth*, Computer Security Resource Center glossary - https://csrc.nist.gov/glossary/term/defense_in_depth↩

CITE THIS ARTICLE

Gurov, P. (2026). *How to put an AI agent on your old laptop with maximum security*. Gurov Digital Research. https://gurovdigital.com/research/ai-agent-security-five-layers

Published at https://gurovdigital.com/research/ai-agent-security-five-layers · Author ORCID 0009-0001-0152-8056