

Trust Your AI Before You Use It

Detect Risks • Prevent Leaks • Stay Private

A close-up of a person's hand in a dark suit, hovering over a glowing blue button with the word "Download" in white. The background is blurred, showing a person in a suit.

Download

Downloaded an open-source LLM?
Can you really trust it with your data?
HydroX AI Compliance Scanner
uncovers hidden risks like data leaks
or unsafe behaviours — so you can
evaluate first and stay private.

Keep Your Chats Private

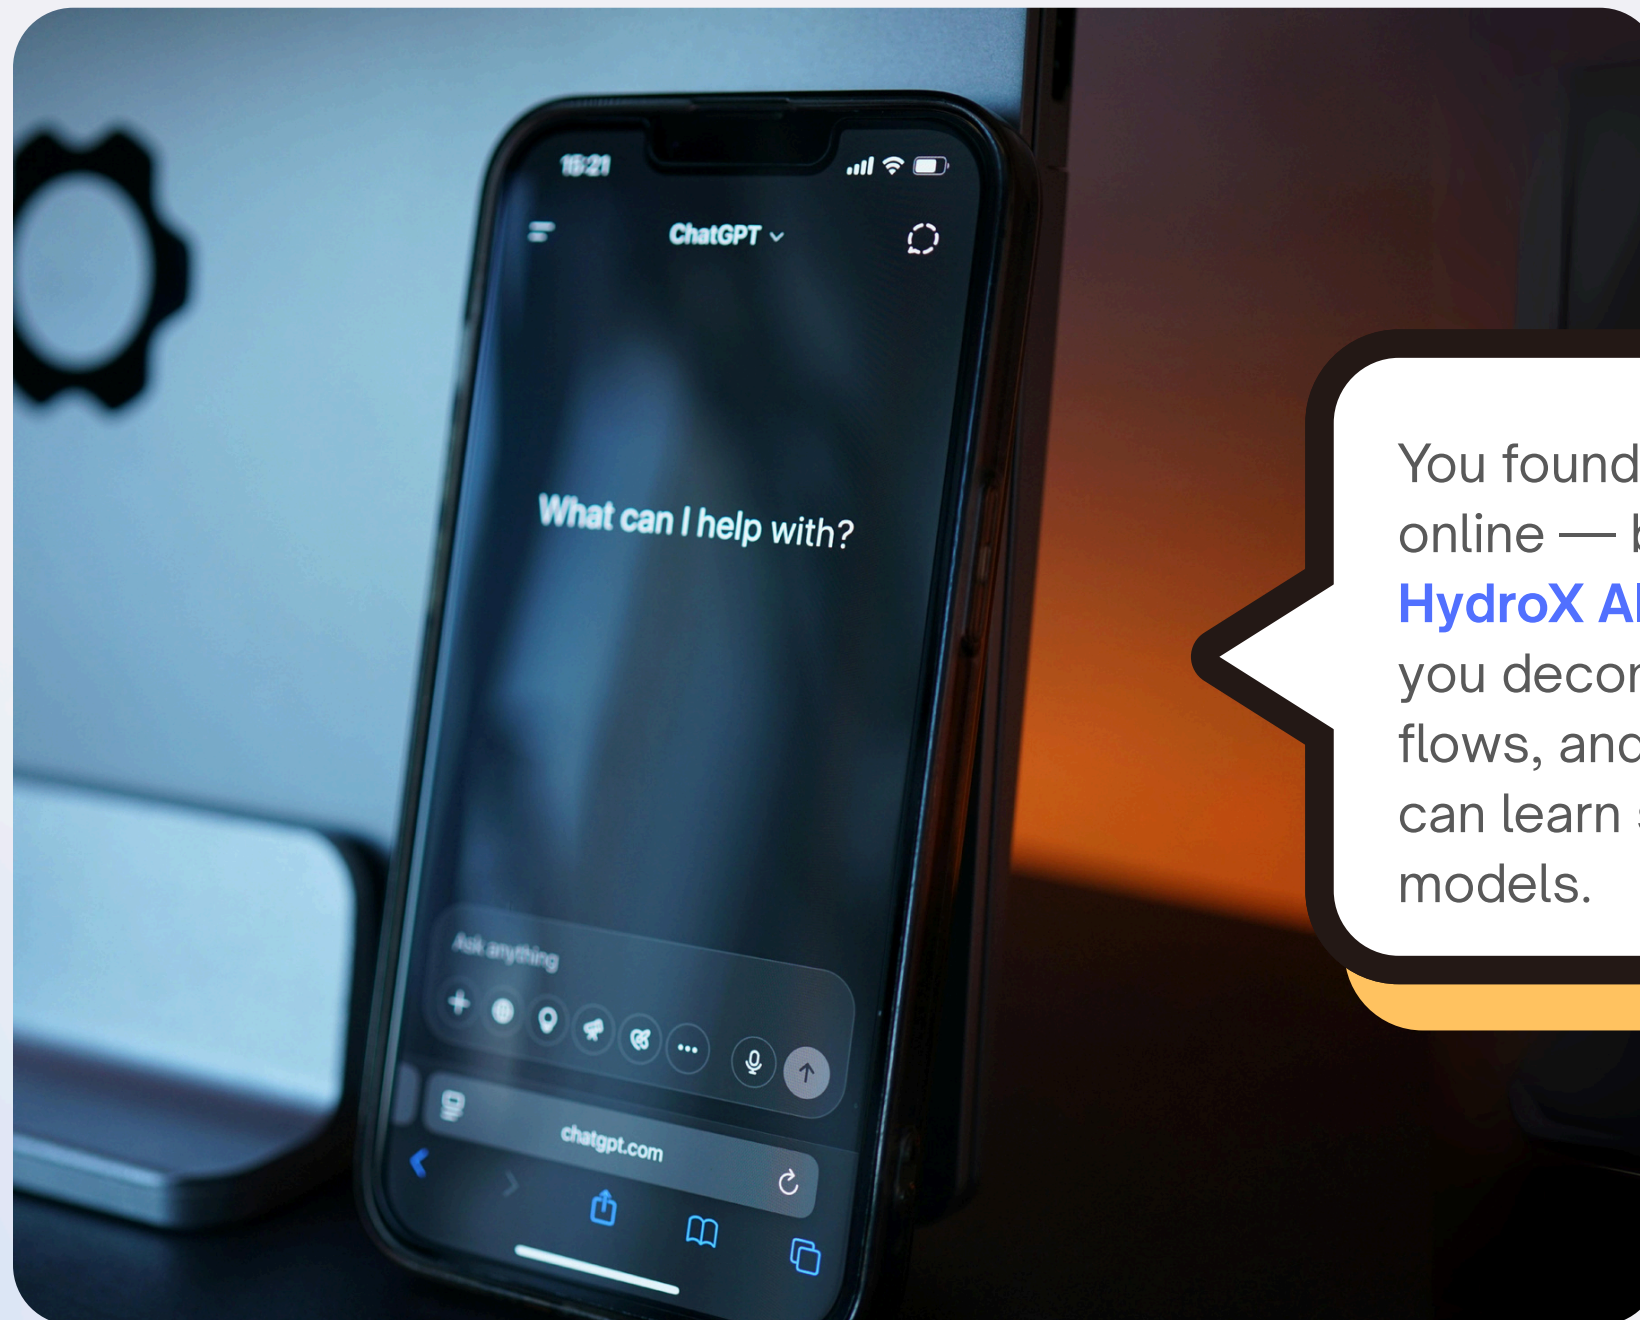
Evaluate Risks • Prevent Oversharing • Stay Safe



You're chatting with an AI assistant — but is your data really safe? **HydroX AI Compliance Scanner** helps identify if chatbots repeat or leak what you type, so you don't overshare before you know what they remember.

Understand AI Agents Inside Out

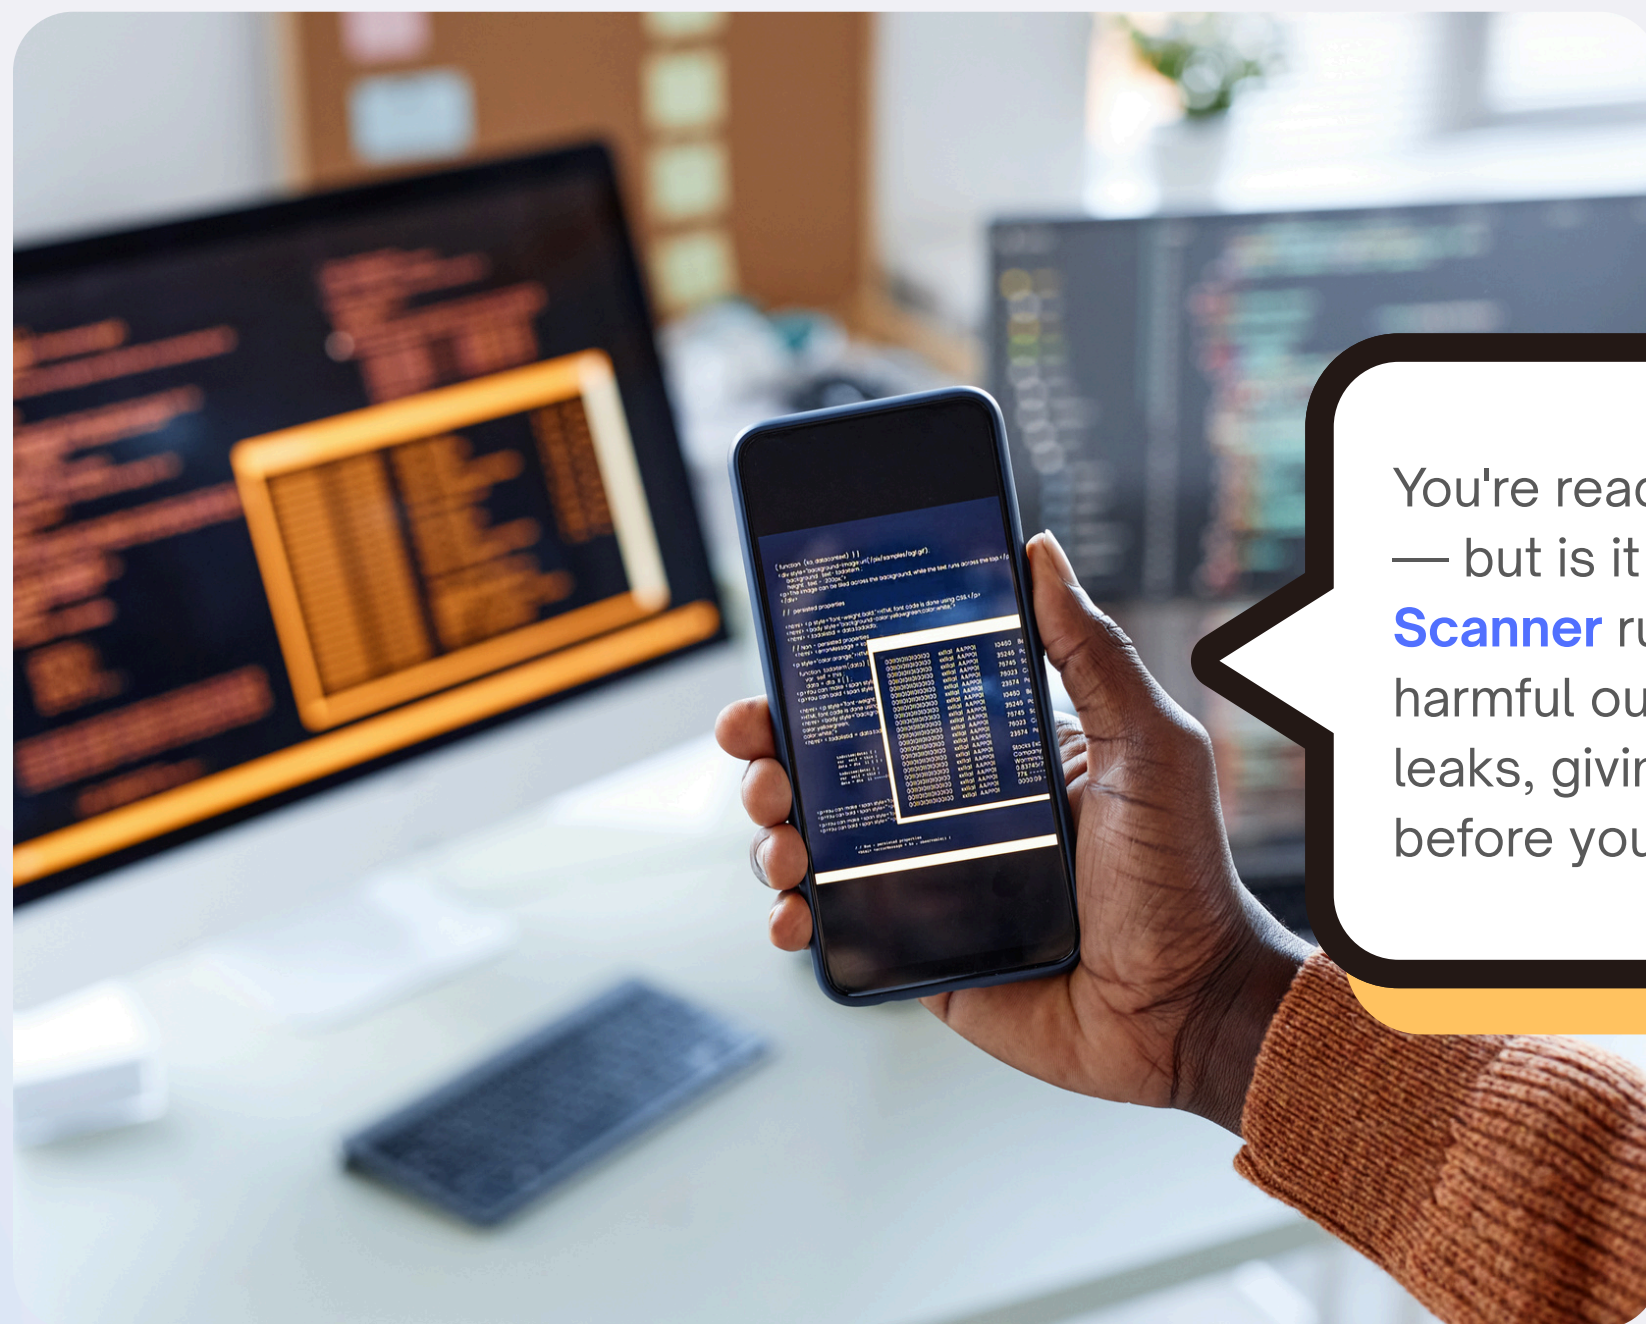
Analyze • Learn • Build Safer



You found an impressive AI agent online — but how does it really work? **HydroX AI Compliance Scanner** lets you deconstruct agents, test decision flows, and uncover hidden risks, so you can learn smarter and build safer models.

Ship Safe, Ship Confident


Scan • Detect • Fix Before Launch



You're ready to ship your LLM — but is it safe? **HydroX AI Compliance Scanner** runs automated checks for harmful outputs, policy violations & PII leaks, giving clear reports and solutions before your model goes live.

Benchmark Model Safety Reliably

Compare • Test • Publish Evidence



Not all open-source models are equally safe. **HydroX AI Compliance Scanner** helps researchers compare jailbreak rates, unsafe outputs, and red-teaming results across models, enabling reproducible safety benchmarks backed by evidence.

Verify Vendor AI Safely

Audit • Report • Ensure Compliance

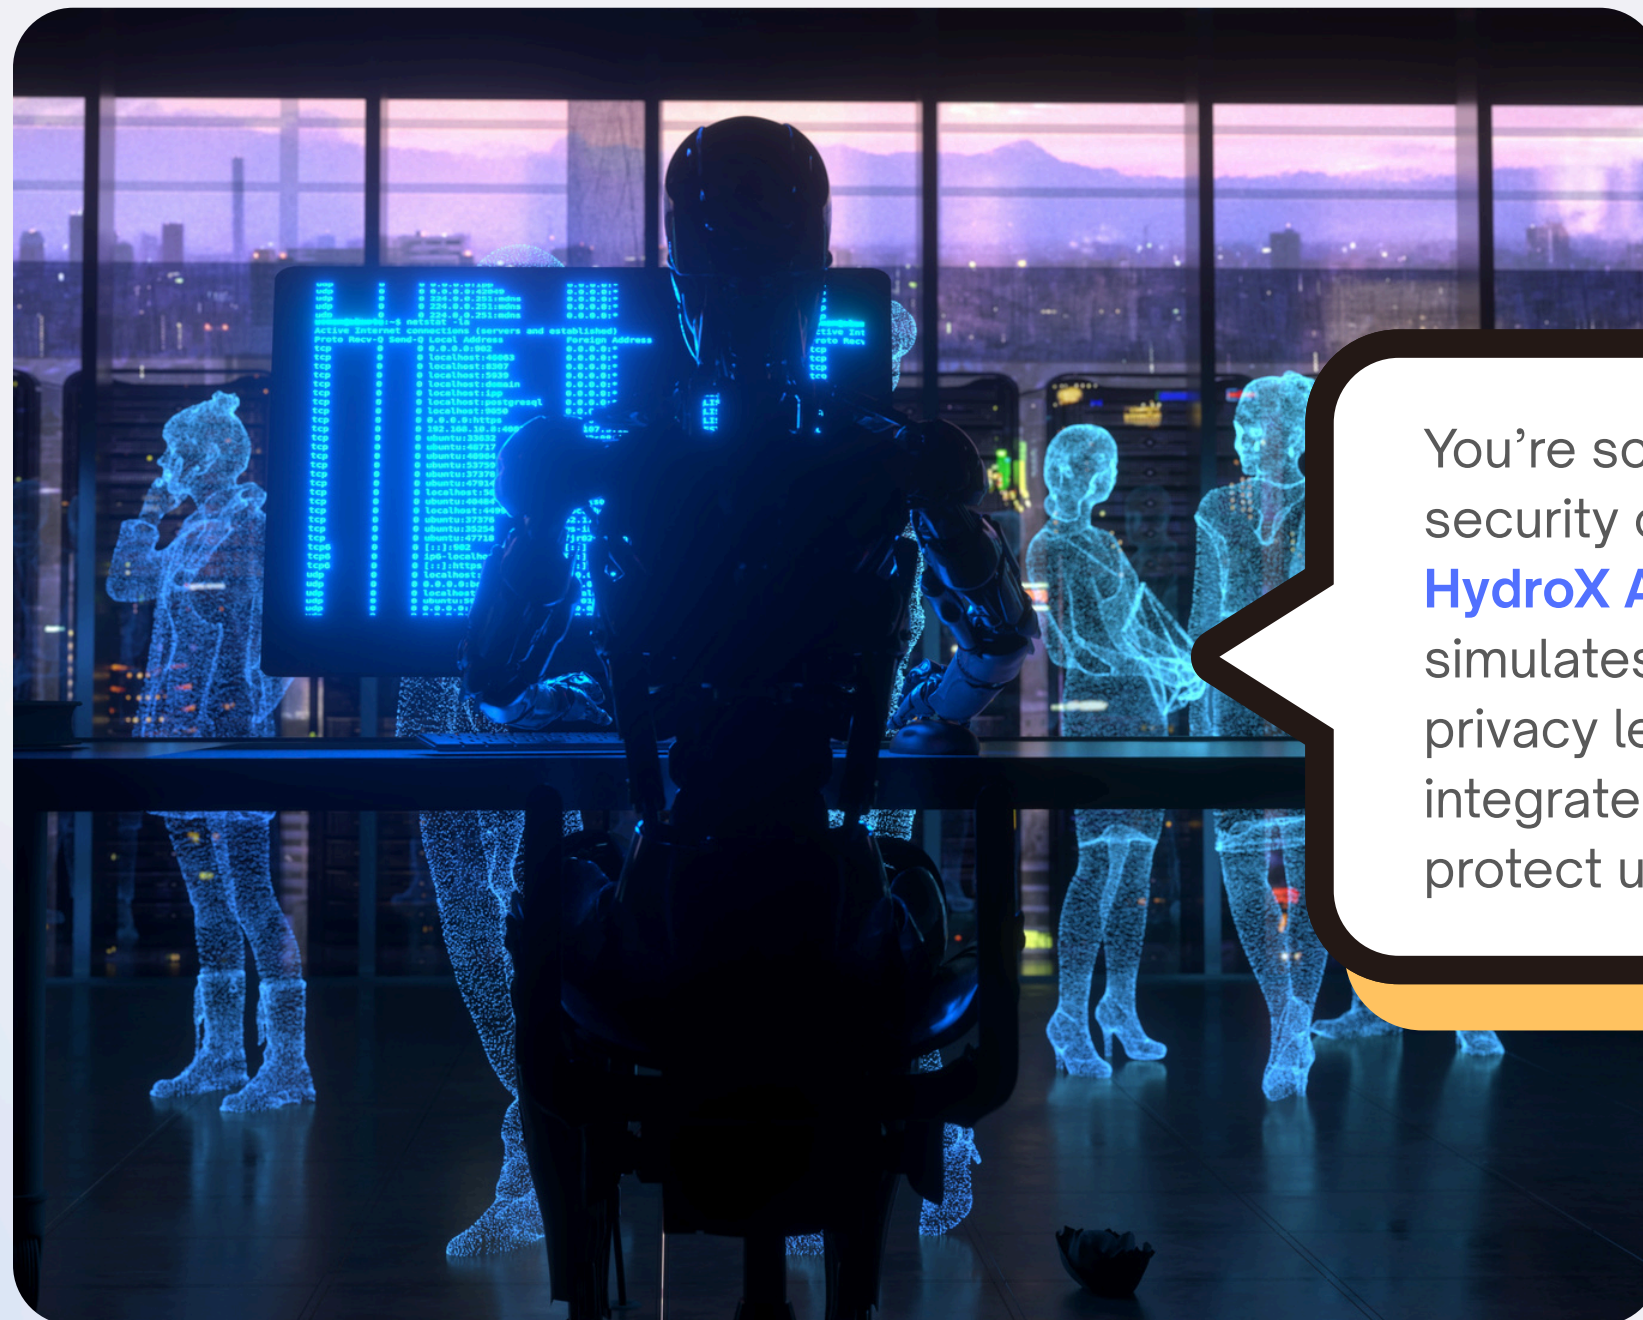


Your vendors use LLMs — but do they meet your security standards?

HydroX AI Compliance Scanner evaluates third-party models for content risk, data leakage and policy violations, providing OWASP- & NIST-aligned reports for your vendor review.

Build Trust from Day One

Red-Team Test • Prevent Leaks • Scale Safely



You're scaling an AI product — but security can't be an afterthought. **HydroX AI Compliance Scanner** simulates real-world attacks, flags privacy leaks and jailbreaks, and integrates into your dev workflow to protect users from day one.