

IN THIS ISSUE

1. Hand work to your bot from X
2. A team of two, tripled
3. Your PC and agents
4. Home AI, no cloud
5. Search photos and audio
6. US money, no equity

The Open Daily

The open AI tools that actually work for your business

A PUBLICATION OF



PubliExpert
PUBLIEXPERT.COM

— FRONT PAGE · AGENTS ON X

You can now hand work to your bot from any post on X: just type @bot

Tag @bot in a reply, a post or a quote, tell it what to do, and the task lands in your Grok Bot. Save, remind, summarize and draft without leaving X.

BY NATHAN ROMANO · PUBLIEXPERT

On Tuesday afternoon, Larsen Cundric, who works on Grok Bot at SpaceXAI, announced on X that @bot tagging is live. The idea is simple: you see a post that matters, reply with “@bot” plus an instruction, and the task goes straight to your Grok Bot. His announcement passed half a million views in under a day.

The examples he gave himself: “@bot add this to my Notion reading list,” “@bot remind me to read this tomorrow,” “@bot summarize this thread” and “@bot draft a reply to this.” It works in replies, posts and quotes. In other words, X stops being just the place where you read things and becomes a place where you hand work to your agent.

Someone already getting a lot out of it is Federico Viticci of MacStories: he uses @bot as the coordinator for several bots, one per internal app on his team, and runs almost all of it from his phone (more below). What the announcement doesn’t say is how much the bot can do on its own once it gets the task; that depends on what you have connected to your Grok Bot.

What does it mean for your business?

X becomes an inbox for your agent. A thread with marketing ideas, a new tool, a competitor launching something: instead of bookmarking it and forgetting it, you send it to the bot with an instruction. One catch: replies on X are public, so whatever you write next to @bot can be read by anyone. No client names, internal prices or personal data on that line.

WHAT YOU CAN ASK IT

SAVE

“add to my reading list”

REMIN

“remind me tomorrow”

SUMMARIZE

“summarize this thread”

DRAFT

“draft a reply”

WHERE

Replies, posts and quotes



ORIGINAL POST
x.com/larsence



VITICCI'S
WORKFLOW
x.com/viticci

NATHAN'S TAKE

Use it to read and organize, not to reply on autopilot. A summary can't embarrass you; a reply posted without a second look can.

Talk to me on X → @nathanromano

EDITOR'S NOTE

Issue No. 3. Today agents step out of the chat window: you can call them from any post on X, and Microsoft wants them working on your PC. We tell you what actually works for you and what isn't worth the price yet. We check everything at the source before it goes to print. Disagree? Tell me on X: @nathanromano.

— NATHAN ROMANO, PUBLIEXPERT

— BRIEFS · HOW IT'S USED

A team of two that tripled its output with @bot and Cursor

Federico Viticci of MacStories laid out his workflow: @bot coordinates several bots, one per internal app (their CDN, article search, proofreading), and each app lives in a cloud Cursor project where agents build it, test it with real data and ask for review. The system watches for its own errors and opens fixes. He has run it mostly from his phone for almost two months, sometimes by voice on walks. “We’re a team of two (plus our designer and dev), but I feel like our output has tripled this year.”



SOURCE
x.com/viticci
TUE, OCT 6, 2026

NATHAN'S TAKE

The lesson isn't the tool, it's the structure: one bot per job, one that coordinates, and everything checked before it reaches you. Copy that even if you use other apps.

Talk to me on X → @nathanromano

AI moves to your desk

— ANALYSIS · YOUR PC

Microsoft wants your PC to be your agents' office. What you get and what you don't

Satya Nadella announced "a new chapter for Windows" on Wednesday: agents working on your PC, securely, on your behalf. The big pieces: MAI-Code-1.1 Flash, a 137-billion-parameter coding model that can run on your own machine; GitHub Copilot handing work off to local models to cut costs; and Copilot taking action directly on your PC with "Hybrid Intelligence."

The fine print, from Microsoft's own pages that day: Copilot's actions on your PC arrive in

the coming months **only on Copilot+ PCs** (with an AI chip of at least 40 TOPS); a powerful graphics card doesn't count. Running MAI-Code at home calls for more than 120 GB of memory: Microsoft only shows it on the 128 GB Surface Laptop Ultra (\$5,899.99; the 24 GB base model is \$2,599.99, ships Oct 16, and isn't officially sold in Mexico). What does work today on any Windows 11 PC: MXC, a sandbox where agents work without touching the rest of your system.



SOURCE
x.com/satyanadella
WED, OCT 7, 2026

NATHAN'S TAKE

Don't buy hardware over this announcement. The same model is already in GitHub Copilot in the cloud, even on the free plan, and that's far cheaper than a six-thousand-dollar laptop. What's worth it today is the sandbox: if you already run agents on your PC, keep them in there.

Talk to me on X → @nathanromano

— NEW · HARDWARE

Oki Home: a home AI computer that keeps your data out of the cloud

Oki Home is an aluminum box that puts your photos, files and messages on one timeline and runs its own 27-billion-parameter AI model around the clock, with no cloud subscription. Inside: an NVIDIA RTX 5060 Ti with 16

GB and 32 GB of RAM; your data lives on a swappable card from 2 to 16 TB. It can also browse sites and handle multi-step tasks. The first batch is \$1,799 with 2 TB, takes a \$59 reservation, and ships in mid-December.



SOURCE
okihome.ai
TUE, OCT 6, 2026

NATHAN'S TAKE

The idea is good: your information, in your house. But it's a new brand's first batch. Wait for the first reviews before you put your business on it.

Talk to me on X → @nathanromano

— OPEN SOURCE · SEARCH

Google releases a model that searches text, photos, video and audio without leaving your device

EmbeddingGemma 2 is open (Apache 2.0, commercial use allowed) and small: 740 million parameters. It understands text, code, images, video and audio together, so you can find a clip with a voice memo. On a phone it needs as little as about 191 MB of memory. The previous version has passed 20 million downloads.



SOURCE
blog.google
TUE, OCT 6, 2026

NATHAN'S TAKE

For your developer: an internal search for your catalog, photos or recordings, with nothing sent to the cloud.

Talk to me on X → @nathanromano

GLOSSARY

Local model

An AI that runs on your own computer instead of in the cloud. Your data stays put and you don't pay per use, but you need a lot of memory: big models call for more than 100 GB.

NUMBER OF THE DAY

120 GB

of memory or more is what Microsoft recommends to run its new coding model on your own machine.

SOURCE: MICROSOFT, OCT 7, 2026

QUOTE OF THE DAY

"I don't want to be dragged down by inefficient tools anymore."

FEDERICO VITICCI, MACSTORIES · ON X, OCT 6, 2026

Have you tagged @bot yet? What was the first thing you asked it? Tell me on X: @nathanromano.

Open source, money & the tip of the day

— OPEN SOURCE · VOICE

OpenOcto: build your own voice assistant that only listens at home

OpenOcto builds voice assistants whose audio never leaves your house: wake word, transcription and voice run on your machine. You pick a personality and a “brain,” cloud or local, so it can work offline. It installs with one command on Mac, Linux or Windows.

Heads up: it's free for personal use, not for your company, and on GitHub it has 9 stars and nothing new since April.



SOURCE
openocto.dev
SEEN OCT 7, 2026

NATHAN'S TAKE

It's a weekend project, not a work tool. Good for understanding how a local assistant fits together; don't put it in front of customers.

Talk to me on X → @nathanromano

TIP OF THE DAY

Use @bot on X without getting burned

Tagging a bot from X saves time, but what you write there is public. Four rules:

- 1 Use it to save, remind and summarize: jobs nobody else sees.
- 2 If you ask for a draft, read it and post it yourself. Never on autopilot.
- 3 Don't put client names, prices or personal data next to @bot.
- 4 Once a week, review what you sent it and clear out what's done.

— BRIEFS · MONEY

The US is again giving startups up to \$2 million with no equity

NSF's Seed Fund has reopened, as Ksenia Moskalkenko pointed out on X: up to \$305K in Phase I and \$1.25M in Phase II (up to \$2M with supplements), with no equity taken. It starts with a 3-page Project Pitch. It re-

quires fewer than 500 employees, a US company doing the work in the US, and at least 50% owned by US citizens or permanent residents.



SOURCE
seedfund.nsf.gov
TUE, OCT 6, 2026

NATHAN'S TAKE

If your company is Mexican, this isn't for you. If you have US partners and build your own technology, the short pitch is worth an afternoon.

Talk to me on X → @nathanromano

PubliExpert AI Guild

Learn to use AI in your business and your life

Join the PubliExpert AI Guild, the community of entrepreneurs building with AI: weekly live workshops on practical AI automation, done-for-you playbooks, templates and a network of people who build with AI, with professional tools included.

guild.publiexpert.com \$67/mo · renews monthly · cancel anytime

Want us to implement it for you? PubliExpert sets it up in your business. Book a call at cita.publiexpert.com



BOOK A
CALL



JOIN THE
GUILD

RECOMMENDED PRODUCT

AFFILIATE LINK



YellowPencil: edit your WordPress design with no code

Click any part of your site and change fonts, colors and sizes visually. It works with any theme, and you can export the CSS whenever you want.

Why I recommend it: agents and search engines read your site; if it runs on WordPress, this lets you make your headings and copy clear without paying for every small change. If you don't use WordPress, it's not for you.

1. envato.market/OJBzNV Affiliate link: I earn a commission if you buy through it. Your price is the same.