



The open voice shell for your
desktop and your AI agents.

always listening

no wake word

any model


open source

Volodymyr Shcherbyna · Ukraine · github.com/vovka/tusk

PROBLEM

Agents can write code and run tools now.

The human side of the loop hasn't changed since 1984.




Slow for everyone. **Painful or impossible** for developers with RSI.

DEMO

Talk, and the desktop does it.


1 • command



"Tusk, open the text editor and type a short poem."

> editor opened, poem typed


2 • pair-coding



"Write a JavaScript function named hello world..."

> function written in your editor

3 • ambient talk



"Anna, did you finish the slides for tomorrow's meeting?"


ignored • not for Tusk



Watch the 90-second demo • tusk.shcherbyna.me

Also by voice: "move it to the left half" • "maximize it" • "set the main agent to gpt-oss-120b"

A cheap model listens, a capable model acts.



Three shifts make this possible now.

01

Sub-second cloud STT and cheap small models arrived

-
- > Always-on listening is affordable

02

MCP became the standard tool protocol


-
- > The desktop is a tool server

03

Coding agents went mainstream

-
- > Developers narrate intent all day

LLM understanding that acts on the real desktop.



Cloud computer-use agents: sandboxed, screenshot-driven, not voice-first.

Working today on Linux. One builder, six months.

[P50 MS]

speech to action, measured

working today

- Always-on voice with gatekeeper
- Dictation
- Model hot-swap

[N]

design partners signed up

- GNOME desktop control
- Voice pair-coding
- Codex backend, tray

[N]

GitHub stars since the public launch

engineering

334 commits · **163** test files

CI on every change

End-to-end tests on a live desktop

Linux/GNOME today · public repo and video: October 2026

Start with developers who need hands-free.

FIRST

Developers who need hands-free

RSI and accessibility users, and keyboard-optional devs pairing with coding agents

Already pay ~\$15/mo for Talon

[N DEVELOPERS WITH RSI · SOURCE]

THEN

Every Linux power user, then Windows and macOS via adapters

[N DESKTOP DEVELOPERS · SOURCE]

LONG TERM

The voice layer other agents plug into

[COMPARABLE: VOICE-TOOL FUNDING · SOURCE]

Open core, with a paid hosted voice stack.

FREE

\$0

Kernel

GNOME adapter

Dictation and coding adapters

CLI and tray

PAID

~\$15/mo

Hosted low-latency voice stack

No API keys to manage

Same price band as Talon

LATER


teams

Windows and macOS adapters

Team licences

Policy and audit

Safety layer first, then Windows and a hosted stack.



TEAM

Built by someone who lives in voice control every day.

Volodymyr Shcherbyna

SOLO FOUNDER

Commercial software since 2007.

Team of one, pair-programming with Claude and Codex.

github.com/vovka

linkedin.com/in/volodymyr-shcherbyna

DevBrother 2021–present

Built ToughCommerce, a multi-tenant B2B concrete-delivery ERP and marketplace (Rails, GraphQL, React, Stripe, QuickBooks, MS Dynamics)

CityFALCON 2018–2020

Financial-news platform at a fintech startup

SoftServe 2013–2017

Core team on Coupa, spend-management SaaS

Education

MSc Computer Science, NTU Kharkiv Polytechnic Institute · Stanford/DeepLearning.AI ML certificate, 2023

Why me: daily Talon voice-control user with custom commands.

Five design partners, and intros to dev-tool angels.

5

Design partners

Linux developers with RSI who use voice daily



Introductions

To dev-tool and accessibility angels

Bootstrapped today, not raising now. A small pre-seed is possible in 2027 for three things: the safety layer, the Windows adapter and the hosted voice stack.

tusk.shcherbyna.me github.com/vovka/tusk volodymyr@shcherbyna.me