NATIVO

COMPANY ONE-PAGER · OCTOBER 2026

Make AI native.

Personal AI that lives with you, not above you. Everything runs on the device: no cloud, no subscription.

THE COMPANY

Nativo builds the OS for on-device AI. **Coquí**, the keychain companion, is the product.
Platform licensing follows once Coquí proves it.
github.com/NativoML · Apache 2.0

MISSION

MAKE AI NATIVE. Offline, private AI/ML that runs on-device.

VISION

Everyone owns the AI they carry. Intelligence that lives with you, not above you.

THE BET

Prove the industry wrong from Puerto Rico. Build the company our engineers don't have to leave.

01 · WHY NOW

Every app ships its own brain
Each AI feature bundles a model plus runtime, 300+ MB duplicated per app, paid in RAM the phone doesn't have.

The platform path is closed
On-device AI exists, behind one company's allowlist, on their terms. OEMs and developers get no open alternative.

Cloud AI is a rental
It dies without signal, bills every month, and your voice lives on somebody else's servers.

AI is leaving the cloud. The only question is who owns what replaces it.

02 · THE STACK

Nativo	**The promise.** The company. Nativo means native to the device. COMPANY
JibarOS	**The soul.** Android 16 fork where AI is a platform primitive. 12 capabilities through one binder AIDL, a kernel for on-device inference. OEMs pick the model per capability. Apache 2.0. OPEN SOURCE
oird	**The engine.** Native inference daemon. Shared model residency: load once, serve many. KV-cache accounting, priority-aware queues. llama.cpp · whisper.cpp · ONNX Runtime · libmtmd. OPEN SOURCE
Governor	**The scheduler.** Sees memory contention and thermals as one coupled loop. Adaptation bounded by QoS contracts. RESEARCH · PROPRIETARY
AgentKit	**The mind.** Commercial intelligence layer above the OS: memory, tools, orchestration. The open platform; the paid brain. COMMERCIAL
Coquí	**The voice.** First proof of the mission: the keychain companion. Standalone-first. Smart without internet. IN DEVELOPMENT



Coquí. The full companion loop, on hardware you own. *Concept rendering; final hardware may differ in appearance and size.*

THE THESIS

SERENO (OSDI '26) schedules memory bandwidth. **ThermE** (arXiv '26) predicts thermal headroom. **Phoenix** (ACM TECS '26) characterizes sustained vision-model thermal behavior. **All three treat symptoms.** On a shared SoC it's one loop: contention burns power, power becomes heat, heat throttles everything, throttling reshapes contention. A scheduler watching one axis misbehaves on the other. The Governor sees both.

`contention → power → heat → throttle → contention`

03 · WHY WE WIN

PRODUCT	WHERE THE AI RUNS	WORKS OFFLINE	SUBSCRIPTION	YOUR DATA
Nativo / Coquí	On the device, in the OS	Yes	None	Yours stays on-device*
Cloud companions	Their servers	No	Monthly	Their cloud
Closed platform APIs	On-device, one vendor's gate	Mostly	Varies	Their terms

* Update checks and health telemetry excepted.

WHO BUILDS IT

Rafael Valle built edge-AI for 100K+ devices at Flock Safety, Android embedded at Motorola, and now leads counter-drone systems as CTO. He knows the surveillance architecture because he built it. Now he's building the opposite: **privacy as architecture, not marketing.**

RAISE

Pre-seed · raising $750K to build. Delaware C-corp in formation · based in Puerto Rico. NSF SBIR pitch in progress as non-dilutive de-risking. 8–10 field-pilot evaluation rigs for air-gapped developers and researchers. No paying customers. The round funds the build; customers come after.