CGEN · INTELLIGENCE REPORT

RYLIS (rylis.io) - private, offline AI

Report ID: IR-DRAFT

Generated: November 08, 2026

BUSINESS

RYLIS (rylis.io) - private, offline AI chat app for IPhone and Android

REPORT

Identity & Overview

Rylis is being developed by Alexander Lavrikov, who is building an application that runs language models directly on a user's phone, so that after a model is downloaded, the local mode can respond without any network connection. The app targets iPhone and Android users. Within the app, users can choose who responds: a local model on the phone, Rylis cloud models, or a third-party service via API. According to the Rylis team, subscriptions are planned only for cloud models and Rylis servers - the on-device mode is intended to function without an active subscription.

The cloud models are divided into three tiers: Standard, High, and Ultra. There is also a "Rylis Private Search" feature.

Publicly findable information about RYLIS's origins, founding date, co-founders beyond Lavrikov, and founding motivation is extremely limited. Lavrikov has stated that he wanted to build a personal assistant that stays available even when connectivity is lost. Rylis is one of several projects run by a team that, according to Lavrikov, has more than ten people and has been working on multiple products for more than two years. No founding date, co-founder names other than Lavrikov, investor history, or specific pivot history is publicly findable as of the date of this report.

MARKET POSITION

1. PocketPal AI

• What They Do: Free, open-source on-device LLM chat app for iOS and Android

• Category: Private mobile AI chat

• Pricing: Free

• Review Score: 3.6/5 (App Store, Google Play)

2. Private LLM

• What They Do: Paid on-device AI chat for iPhone, iPad, and Mac with curated GGUF models

• Category: Private mobile AI chat

• Pricing: Paid - from $4.99 one-time

3. Enclave - Local AI Assistant

• What They Do: Free iPhone app offering offline AI with on-device voice and in-app purchases

• Category: Private mobile AI chat

• Pricing: Freemium

• Review Score: 4.7/5 (App Store)

4. Layla (Layla Network)

• What They Do: On-device AI companion app for iOS and Android, one-time purchase

• Category: Private mobile AI, companion

• Pricing: Paid - $19.99 one-time

• Review Score: 3.8/5 (App Store)

5. Solair AI

• What They Do: Free iPhone offline AI with 60+ models, voice, vision, and health metrics on-device

• Category: Private mobile AI, iOS only

• Pricing: Free

• Review Score: 4.3/5 (App Store)

6. MLC Chat

• What They Do: Open-source mobile LLM runner compiled per device chip, iOS and Android

• Category: Private mobile AI chat

• Pricing: Free

Rylis sits in the on-device / private AI assistant category - specifically the mobile-first, offline-capable LLM chat segment. Offline AI, broadly defined, is AI that runs on a user's own phone or computer rather than on a company's servers, downloading what it needs once and then processing on-device. With offline AI, words and files are processed locally and never sent to a server, so there is nothing to log, nothing to breach, and nothing to use for training. This is the core positioning Rylis occupies.

As of 2026, there are approximately eight to ten local AI apps available for iPhone, including Solair AI, PocketPal AI, Locally (by LM Studio), Private LLM, Enclave AI, Opal AI, MLC Chat, and OneLLM. Rylis is not yet named in those established roundups, reflecting its pre-launch beta status. The key competitors by market segment are:

Traction & Scale

Rylis has not yet launched publicly on the App Store or Google Play. The app is currently in beta and available on iOS via TestFlight. Lavrikov has indicated that independent testing is planned after the app reaches public release on the App Store and Google Play.

No download figures, registered user counts, revenue data, geographic footprint, or third-party growth metrics are publicly available. The team size, per the founding section above, is stated by Lavrikov to be more than ten people. No headcount breakdown by function (engineering, design, marketing) is findable. The developer has stated that the team tested the application on recent-generation iPhones.

The local AI model is downloaded once from Rylis's servers, and the download can run in the background. No app store rank, waitlist size, TestFlight tester count, or beta review data is publicly available. The only independently verifiable milestone is the current TestFlight beta distribution on iOS, with Android status unconfirmed by any source outside the team's own statements.

FINANCIAL PICTURE

No funding rounds, investors, valuation, or revenue figures for RYLIS (rylis.io) are publicly findable in any database, press release, or filing as of November 2026. The company does not appear in Crunchbase, Tracxn, or any startup funding registry found during this research.

The planned revenue model, according to the team's public statements referenced earlier, centres on subscriptions for cloud model access (Standard, High, Ultra tiers). On-device use appears intended to be available without a subscription, though a confirmed public pricing page does not yet exist.

Given the pre-launch stage, there is no basis for estimating annual recurring revenue, monthly active user counts, or lifetime value metrics. Any figure applied here would be invented. The business is best characterised, based on available evidence, as a self-funded or bootstrapped early-stage project, though even that characterisation is unconfirmed.

PUBLIC SENTIMENT

Because Rylis has not publicly launched, there are no App Store ratings, Google Play reviews, Trustpilot entries, G2 listings, or Reddit discussion threads findable for it. The bluescreen.kz article (published November 7, 2026 - one day before this report) is the only independent third-party coverage found, and it is an informational preview rather than a user review.

No social media posts from real users, forum threads on Reddit or Hacker News, or beta tester feedback is publicly findable. The only public voice on the product is Lavrikov himself, whose statements are categorised throughout this report as founder claims, not independent user sentiment.

The absence of public user sentiment is not a negative signal on its own for a pre-launch app; it is simply a data gap. Any characterisation of user reception at this stage would be speculative.

Media & Press

Coverage of RYLIS is extremely thin. The only independently published piece found is a November 7, 2026 article on the Kazakh technology publication bluescreen.kz, which previewed the app's offline capability and described Lavrikov's stated intent. No major English-language technology publications (TechCrunch, The Verge, Wired, Ars Technica), app review sites, podcast appearances, Product Hunt launches, or awards are findable.

No controversies, legal actions, App Store rejections, or significant leadership changes are on record. The app does not yet appear in any "best offline AI" roundup, comparison article, or year-end list in English-language media.

The bluescreen.kz article constitutes the entirety of findable third-party coverage. This is consistent with a product in closed beta that has not yet made a public launch announcement or engaged in formal press outreach.

CURRENT STATUS

Rylis is in active pre-launch development. It is currently in beta and available on iOS through TestFlight. Independent capability testing is planned for after the public App Store and Google Play release. There is no public evidence of a stagnation, pivot, or decline - the trajectory appears to be a team progressing toward a public launch, with no confirmed release date findable.

The competitive environment Rylis is entering is already populated with established free alternatives. PocketPal AI, for example, has 3 million installs on Android alone, representing the scale of the established players. Rylis's differentiation - the combination of local, cloud, and API modes in one interface alongside a "Private Search" feature -

is a plausible but unvalidated product thesis at this stage.

SUMMARY VERDICT

One-line assessment: Rylis is a credible but unproven pre-launch offline AI assistant with a thoughtful hybrid architecture, entering a market already contested by free and established alternatives - the gap between its stated vision and demonstrated traction is currently very large.

Rylis is built around a genuinely useful idea: an AI assistant that works without an internet connection, runs the model directly on a user's phone, and gives users the choice to escalate to cloud models when connectivity is available. The team, according to its own statements, has been building for more than two years and has more than ten people. However, the app is not yet publicly available; it is in iOS TestFlight beta only, with no confirmed Android release date, no public user reviews, no findable funding, and no independent performance data. The founding team beyond Alexander Lavrikov is not publicly named. Every claim about the product's capabilities and roadmap comes from the developer, not from independent testing.

For ordinary people and small businesses, Rylis could become useful as a private, offline AI chat tool - the kind you could use in a location with no signal, or for conversations you do not want sent to any server. Pricing is not publicly listed; on-device use is described by the team as subscription-free, while cloud model access would carry a subscription, but no confirmed price has been published. At this stage the product is not available to the general public, so practical utility is not yet accessible to anyone outside the TestFlight beta.

Sources

• Rylis: ■■-■■■■■■■■■■ ■■■■ iPhone, ■■■■■■■■ ■■■■■■ ■■■■■■■■■■ ■■■■ ■■■■■■■■■■■ - Bluescreen - https://bluescreen.kz/rylis-ii-assistent-dlya-iphone-kotoryy-mozhet-otvechat-bez-interneta/

• 12 Best Offline AI Apps 2026 (Private, No Internet Needed) - https://www.yaps.ai/blog/best-offline-ai-apps-2026

• Best Local AI Apps for iPhone in 2026: All 8 Compared - https://www.solairai.app/best-local-ai-iphone.html

• PocketPal AI for Android - Free App - https://www.appbrain.com/app/pocketpal-ai/com.pocketpalai

FOLLOW-UP ANSWER

Here is the full comparison, structured around the three questions asked.

1. What actually works without internet

Apple Intelligence (iPhone)

The on-device tier handles a defined but limited set of tasks offline. Writing Tools, Genmoji, and Image Playground work offline. For anything heavier - complex Siri queries, image generation at scale, ChatGPT-backed responses - Apple Intelligence analyzes whether a task can be processed on device; when it cannot, the request routes to Private Cloud Compute. So offline operation is real but partial: Apple Intelligence isn't fully offline and wasn't designed to be, though you can use it offline, reliably and daily, for writing, editing, and summarizing.

Google Gemini Nano (Android)

Google built Gemini Nano as its smallest Gemini tier, an on-device model now on version 3 (2026), running fully on-device via Android AICore with zero API cost and sub-second latency. Its offline scope is narrower than Apple's in terms of polish but equally real: it processes many AI tasks directly on device and supports offline features such as smart replies, summarization, and writing assistance. Gemini Nano brings completely offline voice-to-text and text-to-speech to Android 16. However, a flagship device is required; if you are using an older or budget Android

device, you will still be relying on the cloud.

PocketPal AI and similar independent apps

Apps like PocketPal AI and, according to its developer, Rylis work differently from either platform AI. They download a full open-weight language model (GGUF format) to the phone and run inference locally for every query. The entire chat experience - not just specific system tasks - is offline. There is no server tier at all in pure local mode; the model just runs. This is a structurally more complete form of offline capability than either Apple Intelligence or Gemini Nano, but it comes at the cost of model size and generation speed.

2. What stays private - and to whom

The comparison below summarises the structural differences across the three approaches.

| | Sends data off-device? | Who controls the pipeline? |

|---|---|---|

| Apple Intelligence (on-device tier) | No | Apple silicon, Apple OS |

| Apple Intelligence (Private Cloud Compute) | Yes, encrypted | Apple; Apple claims no logging |

| Gemini Nano | No | Google silicon, Android AICore |

| Independent offline apps (PocketPal, Rylis local mode) | No | User's own device; no corporate entity involved |

Apple's privacy architecture is the most publicly documented. Apple's strategy involves on-device processing for sensitive queries, Private Cloud Compute for tasks that exceed on-device capability, and a stated architecture designed so that Apple itself cannot access what users ask their devices. Apple further claims that Private Cloud Compute nodes are specifically constrained from persistent logging, and Apple has committed to publishing the Private Cloud Compute software so that security researchers can verify the claims. These are Apple's claims; independent verification is ongoing, not complete.

Gemini Nano takes a comparable approach on Android. It runs entirely inside Android's AICore service with sub-second latency and full offline operation, with prompts and outputs never leaving the device. Critically, no cloud API or cross-platform equivalent exists; it is exclusively an Android on-device feature.

Independent offline apps offer the strongest structural privacy guarantee, because there is no corporate server in the architecture at all. Running your own local LLM gives full model control, open weights, and absolute offline privacy that no Apple tier matches. This is the gap that apps like Rylis (in its local mode, per the developer's claims) and PocketPal occupy. The trade-off is that you are trusting the app developer not to include analytics or telemetry - a trust claim that, for a pre-launch beta like Rylis, has not yet been verified independently.

3. Is there still room for independent offline AI apps?

There is room, but it is narrowing in one direction and opening in another.

Where the platform AI closes the gap: For the average iPhone or Android user who owns a qualifying flagship and wants writing help, summarization, or voice tasks, Apple Intelligence and Gemini Nano now provide this for free, built in, with no download required. Apple has made "private, on-device AI" a mainstream expectation for roughly a billion device owners - that validation of the local-first premise matters. Any independent app that only matches what the OS now provides for free is competing from a very weak position.

Where the gap remains - four distinct areas:

1. **Device exclusion.** Apple Intelligence needs an A17 Pro chip or newer and at least 8GB of RAM - meaning the standard iPhone 15, iPhone 14, and everything older are excluded entirely. On Android, the flagship requirement for Gemini Nano similarly excludes mid-range and older devices. Independent apps like PocketPal - which has 3 million Android installs per the report - serve users on hardware the platform vendors have locked out.

2. **Model choice and conversation-first interface.** Apple Intelligence and Gemini Nano handle system tasks. Neither is a general-purpose chat interface where you pick a model, run extended conversations, or load a third-party model. Apps like PocketPal and, if it delivers on its stated design, Rylis offer model choice, long-context chat, and the ability to switch between local and cloud models in one interface. That is a different product category, not just a subset.

3. **Absolute privacy for sensitive conversations.** For cloud-tier tasks, Private Cloud Compute extends device-level privacy protections into the data center, though on-device models are necessarily smaller and less capable than frontier AI systems running in traditional cloud environments. For users who want the most capable model *and* a guarantee that nothing reaches a server - ever - the only answer is a fully local independent app. The platform AI systems route complex queries to the cloud by design.

4. **Cross-platform unified experience.** Apple Intelligence is iOS/macOS-only. Gemini Nano is Android-only. An independent app that runs on both platforms with the same interface and model library addresses a cross-device use case neither platform vendor covers.

The honest constraint for Rylis specifically: It is entering this space against established free alternatives (PocketPal, MLC Chat), platform AI that is now built in and free on qualifying hardware, and apps like Enclave (rated 4.7/5) that are already in the App Store. The report notes Rylis is not yet named in any "best offline AI" roundup. Its stated differentiator - combining local, cloud, and API modes in one app alongside a Private Search feature - is a coherent product thesis, but it is unvalidated by any independent test or user review as of this date. The room exists in the market; whether Rylis can occupy it depends entirely on execution that has not yet been publicly demonstrated.

Sources:

• Apple Intelligence: Complete Guide (2026) - https://iphonelifecheck.com/blog/apple-intelligence-guide-2026.html

• Apple Intelligence and privacy on iPhone - Apple Support - https://support.apple.com/guide/iphone/apple-intelligence-and-privacy-iphe3f499e0e/ios

• Is The New Apple Intelligence Feature Actually Usable Offline Or Does It Silently Phone Home - https://www.alibaba.com/product-insights/is-the-new-apple-intelligence-feature-actually-usable-offline-or-does-it-silently-phone-home.html

• Gemini Nano: On-Device Gemini Model Guide 2026 (v3) - https://hokai.io/hub/models/gemini-nano

• Gemini Nano 2026: Powerful On-Device AI, Features, Benefits, and Everything You Need to Know - https://www.online-sciences.com/technology/gemini-nano-2026

• Offline Voice AI on Android: Gemini Nano Voice Explained - https://www.freevoicereader.com/blog/android-gemini-nano-offline-voice-ai

• Apple WWDC 2026: On-Device AI and the Privacy Bet - https://www.vaasblock.com/news/apple-wwdc-2026-on-device-ai-strategy-privacy-bet/

• Apple WWDC 2026: On-Device AI vs Local LLMs Analysis - https://www.promptquorum.com/local-llms/apple-on-device-ai-vs-local-llms