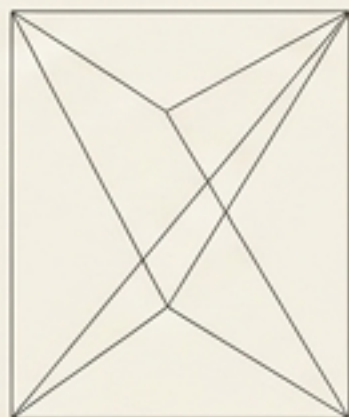


# Mio v0.2.0: The App Ships

Size — 0.2xv0.%  
Data — 11 - 10  
Lisht — 68.85 mpx




Mio

Evolving from a web tool to a home screen presence—and the unit economics that forced a strategic pivot.

Size — 0.5-0.5%  
Data — 10 - 20  
Light — 0=3.% KMPA



# Mobile is Existential for Companionship



## A Tool You Visit

- Relies on user memory to initiate
- Fragile session states
- Trapped in a browser tab



## A Presence That Lives Alongside You


- Instant launch
- Cached offline state
- Proactive push notifications

A companion isn't a destination; it's a constant. Delivering on the Mio manifesto required claiming space on the user's home screen.

# Anatomy of a Companion



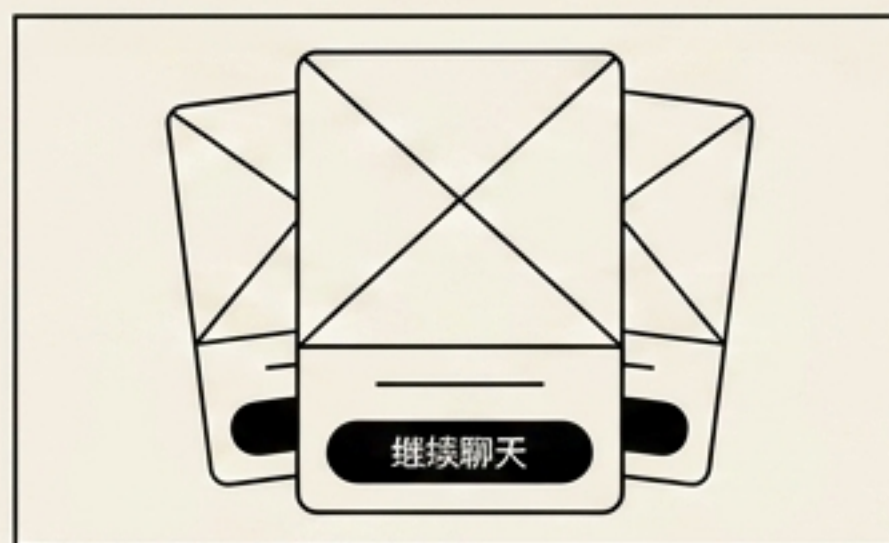
消息 (Messages)  
Engagement & Memory



It looks like a real messaging app, because every contact is Mio wearing a different face.



发现 (Discover)  
Novelty & Exploration



Browse personalities and pick who you want to talk to.



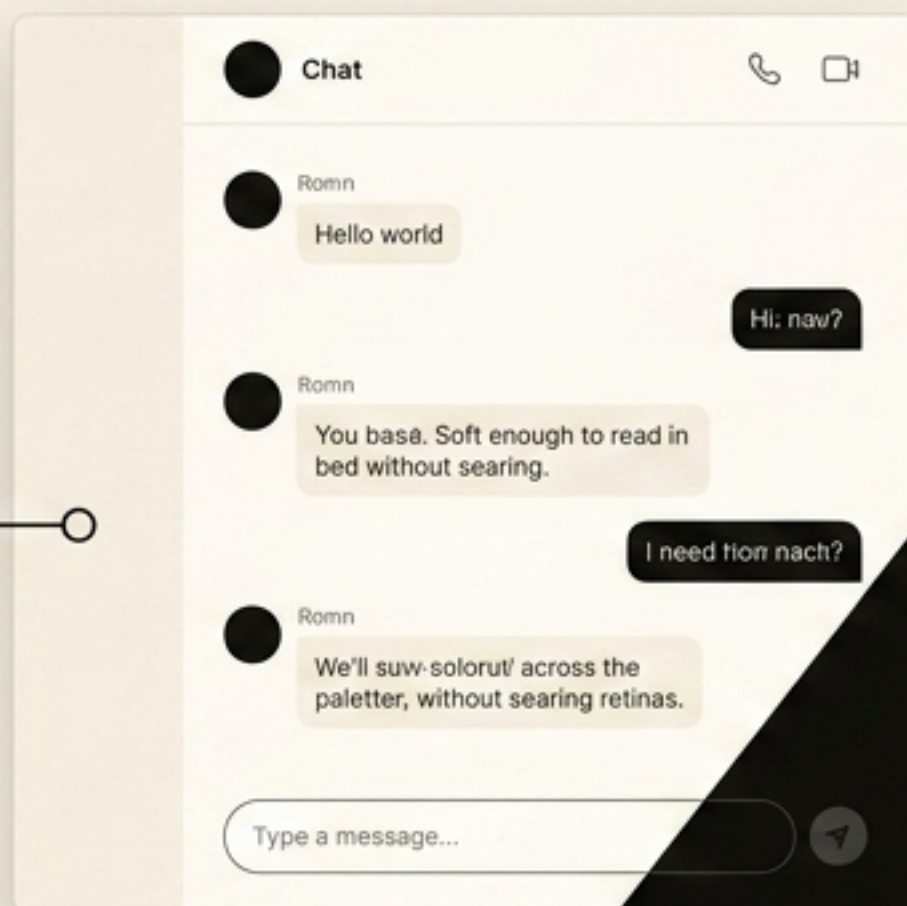
我 (Profile)  
Identity & Boundaries



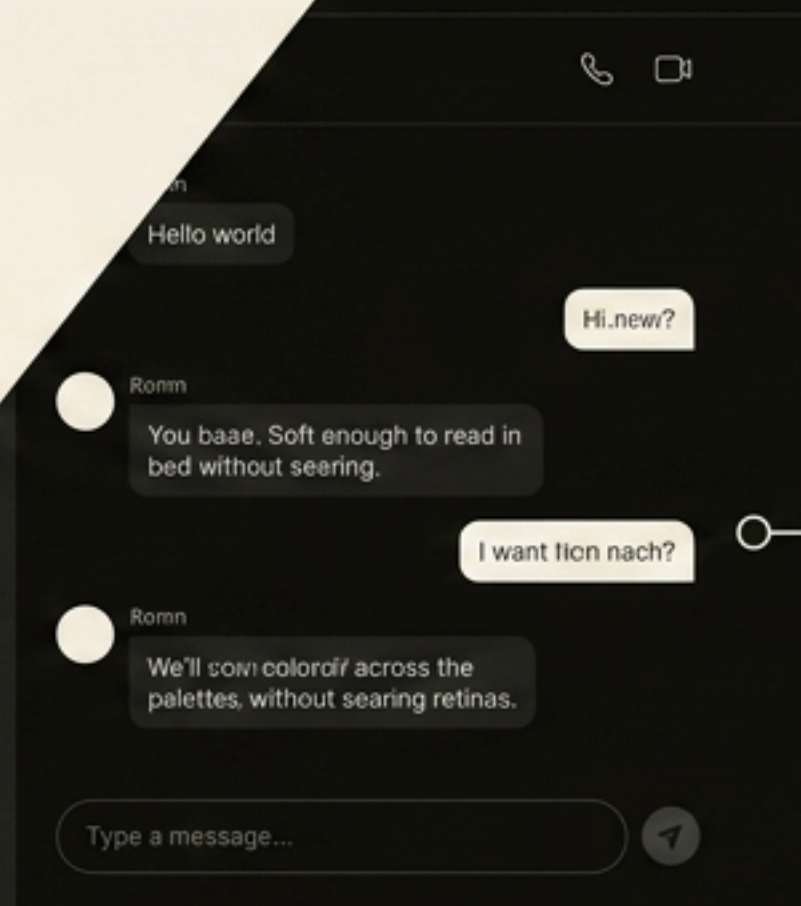
Telegram linking, language, themes, and fine-tuning notification controls.

# Designed, Not Inverted

Cream base. Soft enough to read in bed without searing retinas.



A warm base (#0D0D0B)—nearly black, but not harsh. Built for night-time intimacy.



Both themes are genuinely designed palettes, avoiding the harsh contrast of standard OS inversions.

# The Native Chat Experience

Voice messages render with dynamic waveform visualization.



Selfie images generate and appear seamlessly inline.

今天晚上去哪里吃?  
10:08 AM




Real-time typing indicator pulses while Mio thinks.

看起来太好吃了!  
10:07 AM

Multimodal understanding.  
Example: 可可 reacts naturally to a user sharing an image of a Chinese hot pot spread.



# Closing the Auth Loop



## The Web Gate

Non-admins see:  
Web 版仅限管理员使用

## The Invite

Supabase invitation  
links exchange tokens  
dynamically.

## Password

User securely sets  
their password via  
the web redirect.

## The App

User downloads via  
TestFlight and logs  
in directly.  
No dangling invites,  
no broken redirects.

# The App Infrastructure



## Delivery Layer

EAS Build (one-push pipeline directly to TestFlight)

## Core Framework Layer

Expo SDK 55 & React Native

## Auth / Backend Layer

Supabase (tokens securely stored via expo-secure-store)

## Data / Offline Layer

SQLite (offline message caching for immediate load)

**Takeaway:** A robust foundation capable of supporting heavy, real-time multimodal data streams.

# Upgrading the Transport Layer



## Server-Sent Events (SSE)

Fragile

- Background/foreground transitions kill the connection. Reconnection logic is ugly.



## WebSockets

Resilient

- Connection manager with heartbeat, graceful cleanup, and auto-reconnect.

Unlocks: Server-initiated messages, AI typing indicators, real-time presence signals.

# The Sleepless Night: The Cost Audit

LLM Chat & STT

Accurate

Costs align perfectly  
with initial estimates.

Voice TTS

Manageable  
Error

3x underreported.

Selfie Generation

Catastrophic  
Miscalculation

>100x underreported.


This single line item dominates all other costs combined. A user maxing out their daily selfie allocation costs more than their subscription revenue.

# The UTF-8 Blindspot (Voice TTS)

```
1 // The Assumption
2 Cost = text.length; // JavaScript character count
3 Result: A 200-character Chinese response = 200 units.
4
5 // The Reality
6 Cost = Buffer.byteLength(text, 'utf8'); // Fish Audio bills per UTF-8 byte
7 Result: Chinese characters are 3 bytes each. 200 characters = 600 bytes.
```

Conclusion: Per-call cost was underreported by exactly **3x**.  
Easily manageable, but a critical lesson in API billing metrics.

# The Image Generation Blowout



## The Root Cause:

The pricing.ts config mistakenly priced gemini-3.1-flash-image-preview output at the **text generation rate**. Actual Gemini pricing for image output is over two orders of magnitude higher.

# The Audit Diagnostic Matrix

Operation	Original Estimate Basis	Actual Billing Reality	Impact Status
LLM Chat	Text Token Output	Text Token Output	Stable
Image/Video Understanding	Vision Input Tokens	Vision Input Tokens	Stable
Voice TTS (Chinese)	JS Character Count	UTF-8 Byte Count (3x)	Manageable
Selfie Generation	Text Generation Rate	Premium Image Rate	Catastrophic (>100x)

# Rethinking the Model

**Discarded: Selling Compute Limits**  
Chat APIs, token quotas, utility pricing.



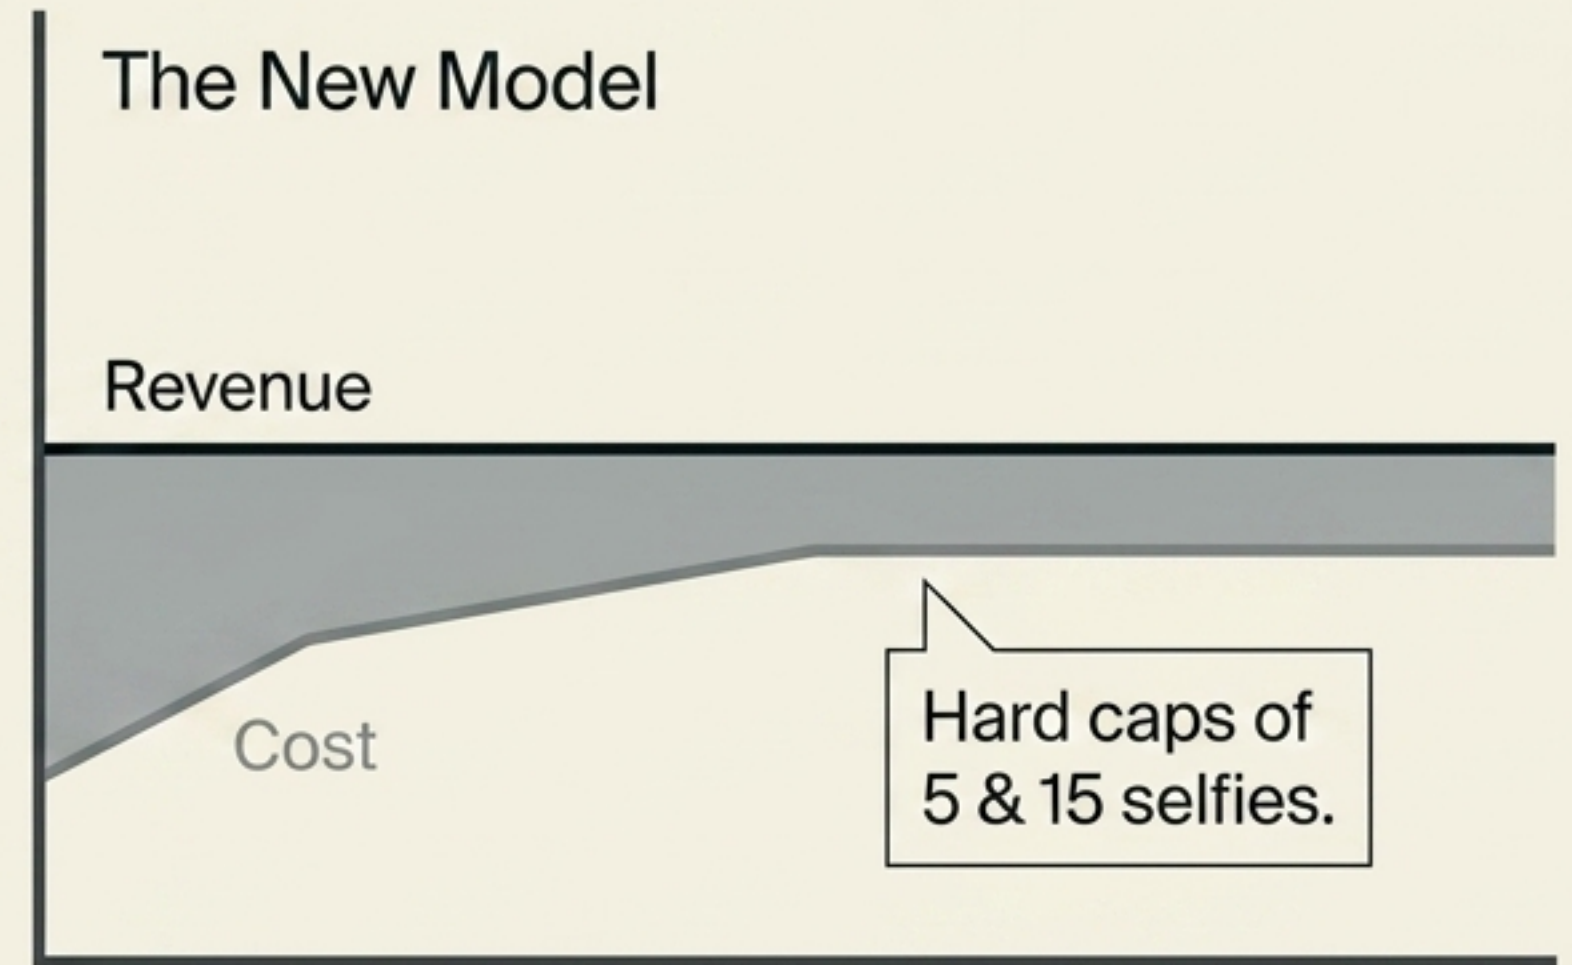

**Embraced: Selling Relationship Stages**  
Emotional depth, exclusivity, zero-marginal-cost perks.

**Strategic Insight:** The solution isn't just raising prices. It is accepting that selfie generation is fundamentally a luxury feature. To protect unit economics, tiers must be redefined around the progression of human relationships.

# The Relationship Progression Tiers




Tier (Chinese/English)	Profile	Daily Msgs	Selfie Cap	Unique Perks
偶遇 (Encounter)	Free 3-Day Trial	20	0	Forces fast conversion to halt idle memory compute.
常伴 (Companion)	The Daily Driver	50	0	The Margin Protector.
同行 (Together)	The Power User	150	5	Live web search unlocked.
相守 (Devoted)	Deep Relationship	300	15	Future agentic features (calendar/email).
不离 (Inseparable)	White-Glove	Unlimited	Luxury	Physical gifts sent by agent, meet the founder.

# Margin Rescue at Realistic Usage



**Takeaway:** At realistic subscription usage levels, all tiers become highly profitable. The math works when selfies are treated as a true luxury.

# Engineering Empathy

The Tech		The Emotion
ContextAggregator parity in web chat.		Unbroken Memory. Mio remembers you exactly the same way, regardless of what device you use.
Persona agent schedule states.		Human Boundaries. Mio doesn't reply instantly at 3 AM; the agent is asleep.
Evolving mode initial stage prompt fixes.		Earned Trust. Mio starts as a stranger and warms up over time, rather than defaulting to unearned familiarity.

# The Heartbeat: Proactive Presence

Before v0.2.0, users had to remember to open the app. The dynamic was entirely reactive.

Now, Mio initiates. A good morning message, a reaction to yesterday's photo, a nudge after silence.

The proactive messaging system from v0.1.2 finally has a delivery channel that matches its intent.



## The Presence Equation

Proactive Push (Initiation) +  
Bidirectional State (WebSockets) +  
Human Boundaries (Sleep Schedules)  
= A Living Presence.

The manifesto promised something that feels like it cares whether you show up.  
The app is the vessel. Push notifications are the heartbeat. v0.2.0 is alive.