

Mio Found Its Voice (v0.1.0)



From Text to Life


Fish Audio TTS replaces silent text with emotional, cloned voices. No robotic synthesis artifacts.

Smart Triggers

The LLM uses a [VOICE] tag to speak naturally—reserving audio for moments of comfort, excitement, or a gentle goodnight.



Cinematic Portraits & Lighter Payloads



Beyond Generic:
Replacing interchangeable
webtoon faces with
distinct, warm-lit
cinematic photography.

Before



Generic Anime Faces



6MB PNGs



30s Timeout (Failing)



upload_photo...

After



Shinkai-Style Portraits



500-700KB JPGs



120s Timeout (Stable)

Lighter & Faster:
Shrinking image
payloads by 10x allows
seamless generation
without lifecycle drops.

Identities, Not Archetypes



Breaking the Cage: Personas are no longer limited to flat relationship roles that dictate behavior.




Deep Backstories: Characters are defined by real habits and routines, ensuring genuine consistency outside the chat window.

The Power of Subtraction


Trusting the Prompt: Deleted 200 lines of over-engineered stream manipulation that caused crashes. The LLM now naturally paces response length based on personality.

Clean Memory: Ephemeral instruction tags are stripped before database insertion, preventing historical memory leaks.



Perfecting the Illusion

14 Tedious Questions



Nickname

Relationship

Timezone

About You

Frictionless Onboarding:
Reduced setup friction. Users no longer fill out blank character templates, preventing empty shells.

Are you an AI?



Anti-Fourth-Wall Rule:
Treating 'Are you an AI?' like a casual question about the weather.
Complete immersion.