

Sesame Raised \$300 Million. What Has Actually Been Shipped?

An outside-in investigation of Sesame AI's public product surface, funding, open model, beta app, and what actually exists beyond the demo.

2026.03.23 01:20 DENNIS HEDEGREEN INVESTIGATE OPEN-NOTE

[HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/SESAME-WHAT-HAS-ACTUALLY-BEEN-SHIPPED/](https://hedegreenresearch.com/articles/sesame-what-has-actually-been-shipped/)

Sesame is on every breakfast table in Denmark.

You do not think about it.

You just eat.

That is roughly how most people have consumed Sesame AI so far.

A demo appeared.

It sounded incredible.

People talked about it for a week.

Then they moved on and forgot to ask the obvious questions.

This is an attempt to ask them.

Important:

this article evaluates Sesame AI from the outside.

Where a claim is based on public reporting or directly observable public material, it is treated as source-backed.

Where I describe my own use of the demo, it is marked as firsthand observation.

Where I draw a conclusion from the visible surface rather than documented internal fact, it is treated as inference.

What Is Clearly Real

Some things are not in dispute.

Sesame AI exists.

It was founded in June 2023 by Brendan Iribe and Ankit Kumar.

Iribe co-founded Oculus VR and led it through the two billion dollar Facebook acquisition.

Kumar built AI systems at Ubiquity6 and later led engineering on Discord's Clyde AI.

[Source: TechCrunch, October 2025; Contrary Research company profile; Andreessen Horowitz investment announcement, February 2025.]

The company has shown unusually strong voice technology.

Its demo made an immediate impression because it sounded less like a voice assistant and more like an actual conversational presence.

The pauses landed better.

The breathing sounded more natural.

The rhythm felt less mechanical than what most people were used to hearing.

Voice AI has had a persistent problem for years.

Even when it works, it usually sounds dead on arrival.

Sesame appears to have pushed further through that barrier than most of its competitors, at least in demo form.

[Firsthand observation, March 23, 2026]

I used the demo at `app.sesame.com`.

Each session ended after approximately 30 minutes.

A new session could be started immediately.

The system retained conversational context within the session.

When I mentioned `hedegreenresearch.com` without pasting a direct URL, the AI later referenced the site in a way that suggested some form of retrieval or lookup.

I did not measure latency directly.

I also cannot determine from this alone whether the observed retrieval behavior reflects persistent memory, within-session context handling, or a stable search feature.

What I can say is simpler:

the interaction felt unusually fast and unusually natural.

That is one observation, not a product audit.

[Observed by others]

Sequoia Capital reported that more than one million people used the demo in its first weeks, generating more than five million minutes of conversation.

Multiple reviewers described the voice quality as the most realistic they had encountered.

Sesame's own co-founder has said the system is still too eager and often inappropriate in tone and pacing.

[Source: Sequoia Capital partner blog, October 2025; Beebom, TechRadar, PCWorld, Ars Technica, March 2025; Sesame research page.]

The team credibility is also real.

This is not an anonymous startup with a landing page and a promise.

Nate Mitchell, another Oculus co-founder, joined as Chief Product Officer in June 2025.

Hans Hartmann, former Oculus COO, serves as COO.

Ryan Brown spent a decade between Oculus and Meta Reality Labs.

[Source: Contrary Research; Sequoia Capital blog; TechCrunch, October 2025.]

So the point of this piece is not to pretend Sesame is fake.

The point is narrower.

What has actually been delivered?

What Has Actually Been Shipped

As of March 2026, the public inventory is short.

Two things are clearly real on the public side.

First.

A web-based voice demo, launched February 27, 2025.

It is free.

It is labeled a research preview.

It features two voice characters named Maya and Miles.

In my own use on March 23, 2026, sessions ended after approximately 30 minutes, but a new session could be started immediately.

The demo has been updated since launch, but the scope and nature of those updates are not fully documented from public sources.

[Source: Sesame company website; GIGAZINE, March 2025; firsthand observation.]

Second.

An open-source speech model called CSM-1B, released March 13, 2025 under an Apache 2.0 license.

It is a one billion parameter base model built on a Llama backbone with a specialized audio decoder.

Sesame trained three sizes: 1B, 3B, and 8B.

Only the smallest was released publicly.

The fine-tuned Maya and Miles voices were not included.

No training code was provided.

[Source: GitHub repository SesameAILabs/csm; Hugging Face model card; TechCrunch, March 2025.]

The model has community traction.

The public GitHub repository shows substantial stars and forks.

The Hugging Face page shows significant usage.

But the public development surface has looked quiet for a long time.

As checked on March 23, 2026, the repository showed 24 total commits and no new commits after May 27, 2025.

Users on Hacker News also noted that inference on an RTX 4090 takes 5 to 10 seconds per sentence, far from the demo's reported response time, and criticized the release as a much weaker public version than the system people actually heard in the demo.

[Source: GitHub repository and commit history; Hugging Face model page; Hacker News discussion threads, March 2025.]

Beyond those two things, the surface becomes thin.

An iOS app was announced alongside the Series B in October 2025.

Public reporting says it is distributed through TestFlight, requires invitation access, and includes confidentiality restrictions.

As of March 23, 2026, the Sesame website still showed a beta waitlist rather than a normal public product page.

No public Android version was visible on the company surface I checked.

[Source: TechCrunch, October 2025; Apple TestFlight listing; Sesame website as checked March 23, 2026.]

Smart glasses were also described as part of the company vision.

Lightweight.

Audio-only.

No AR display.

[Source: Sequoia Capital blog, October 2025; SiliconANGLE, October 2025.]

But as checked on March 23, 2026, the public surface remained thin.

The Sesame website and press coverage reviewed for this article contained no hardware specifications, no named manufacturing partner, and no shipping date.

A search of the FCC equipment authorization database also returned no regulatory filings under Sesame or obvious related terms.

That does not prove these things do not exist internally.

It means they were not visible on the public surface I reviewed.

That is the full public inventory this article is willing to stand on:
a demo, a base model, a waitlist, and a concept image.

The Money

Sesame has raised over 300 million dollars across three rounds in under two years.

Crunchbase reports 307.6 million.

PitchBook reports 322 million, possibly including undisclosed capital.

This article therefore uses the more conservative over \$300 million framing.

[Source: Crunchbase; PitchBook; TechCrunch, October 2025; Bloomberg, March 2025.]

Public reporting describes:

- a seed round of roughly 10 million
- a Series A of 47.5 million led by Andreessen Horowitz
- a Series B of 250 million co-led by Sequoia Capital and Spark Capital

[Source: Contrary Research; Andreessen Horowitz investment announcement, February 2025; TechCrunch, October 2025; Bloomberg, March 2025.]

Total visible revenue:

none that I have been able to verify from the public surface checked for this piece.

As of March 23, 2026, a review of the Sesame website, Apple App Store, Google Play Store, and major press coverage returned no pricing page, no paid tier, no API offering, no enterprise plan, and no subscription product.

That does not prove revenue does not exist.

It means no commercial offering was visible on the public surface.

Companies build in private for legitimate reasons.

But over 300 million dollars and more than a year after the viral moment, the gap between capital raised and product shipped is still wide.

There is one personnel detail worth noting.

Johan Schalkwyk, listed as the ML lead on the core CSM research, later left Sesame for Meta's Super Intelligence Lab.

A single departure does not by itself indicate trouble.

But in a company of roughly 40 to 60 people where the core technology is the primary asset, it belongs in the visible record.

[Source: Contrary Research; R&D World, March 2025.]

The Competitive Problem

The rest of the market is not waiting.

OpenAI sells access to voice through a Realtime API and consumer subscriptions.

ElevenLabs operates a commercial voice platform with paid tiers and enterprise sales.

Google offers Gemini Live as part of its public consumer AI products.

Meta sells Ray-Ban smart glasses with built-in voice AI that a customer can purchase today.

[Source: OpenAI product pages; ElevenLabs website; Google product pages; Meta product listings.]

These companies all have public products in market.

Sesame, from the visible outside, still does not.

If Sesame's main edge is conversational naturalness, that edge holds only as long as it remains distinctive.

And competitors with shipping products are still improving.

The question is whether Sesame can convert a remarkable demo into a real product before the window narrows.

That is an execution question, not a moral one.

What the User Might Be Paying With

This section is inference.

I am raising a question, not reporting a documented fact about Sesame's internal data practices.

The demo is free.

The sessions are substantial.

The AI listens to how you talk, what you talk about, and how you follow up.

Sesame has publicly highlighted the scale of early conversational usage.

Over five million minutes were reported in the first few weeks alone.

[Source: Sequoia Capital blog; firsthand observation.]

A company trying to build the most human-sounding voice AI in the world is offering free access to a system that collects exactly the kind of speech patterns such a model would find valuable.

I have not reviewed Sesame's privacy policy for this piece, and that is a real gap I am naming openly.

But the structural alignment between free product and training incentive is still worth noting.

This is an open question, not a concluded finding.

The Strongest Counterargument

The strongest defense of Sesame is simple.

Hardware takes time.

Good voice systems take time.

The team built Oculus, which also looked like a demo long before it became a two billion dollar acquisition.

That is fair.

And if Sesame eventually ships strong glasses and a real consumer product, the current gap may later look like a normal build phase.

That possibility should stay open.

But it does not erase the present-tense question.

What Can Be Said Honestly

Sesame has remarkable voice technology.

Real team credibility.

Serious investor conviction.

But from the public side, the shipped surface is still a free demo, a quiet open-source surface, a locked beta, and a concept image.

That is a narrow foundation for a billion-dollar story.

That may change fast.

These people have done it before.

But as of March 2026, the voice is ahead of the product.

And until more of that surface becomes real, Sesame should be understood not as proof, but as a bet.

A very expensive, very well-spoken bet.

Follow the data, not the narrative.

— Dennis Hedegreen, follow the data

Method Note

This article was built from the following source categories:

Company primary sources:

Sesame website, Sesame research page, GitHub repository (SesameAILabs/csm), Hugging Face model page (sesame/csm-1b), Apple TestFlight listing, visible public beta surface.

Investor primary sources:

Andreessen Horowitz investment announcement, Sequoia Capital partner blog.

Press reporting:

TechCrunch, Bloomberg, SiliconANGLE, R&D World, GIGAZINE.

Research aggregators:

Contrary Research, Crunchbase, PitchBook.

Community sources:

Hacker News discussion threads, Hugging Face community data, GitHub commit history.

User reports:

Beebom, TechRadar, PCWorld.

Firsthand observation:

author's direct use of the Sesame web demo on March 23, 2026.

Absence checks performed on March 23, 2026:

Sesame website, Apple App Store, Google Play Store, major tech press, and the FCC equipment authorization database using obvious Sesame-related search terms.

Where the text states a claim as fact, a source is noted inline.

Where the text describes direct experience, it is marked as firsthand observation.

Where the text draws a conclusion from available evidence rather than documented internal fact, it is treated as inference.

— Dennis Hedegreen, follow the data