

CONSUMER AI BRIEFING

The State of Voice Assistants

Vol. 1 · July 2026

PLATFORM ASSISTANTS: THE GREAT LLM SWAP

CHATGPT & NEW ENTRANTS

VOICE MODELS & TECH

ENTERPRISE VOICE AGENTS

TRUST, SAFETY & THE FRAUD WAVE

USAGE & MARKET



BY DANIEL SHANKLIN

danielshanklin.com/stateof/voice

Executive Overview

The quarter voice grew up. Every major platform swapped its legacy assistant for an LLM brain, OpenAI made conversation full-duplex, enterprise voice agents raised ~\$1.7B in eleven weeks — and voice-clone fraud went industrial. The interface is finally natural; so is the crime.

900M

ChatGPT weekly active users (Feb 27) — the distribution engine behind GPT-Live voice

[TECHCRUNCH](#)^[1]

~\$1.71B

Enterprise voice-agent funding across 9 disclosed rounds in ~11 weeks (per-outlet reports)

PER CNBC / TECHCRUNCH / FORTUNE, COMPILED

~\$900M

US losses to AI voice-cloning scams, FBI figure (as reported June 2026)

PER MALWAREBYTES / MONEYWISE

~\$1B/yr

Reported value of Apple’s multi-year deal for Gemini models to power the rebuilt Siri (Bloomberg est.; unconfirmed by either company)

[REUTERS](#)^[2]

\$99.99

Google Home Speaker — first new Google speaker in six years; Gemini formally retires Assistant

[BLOG.GOOGLE](#)^[3]

7M+

Meta AI glasses units sold in 2025 — tripled YoY (EssilorLuxottica)

[CNBC](#)^[4]

+1,600%

Reported rise in voice deepfake attacks (as reported May 2026)

PER HACKERNOON

600M

Alexa devices in the installed base (Panay claim, Feb 2025) — no official Alexa+ subscriber count since “1M+” in Jun 2025

[CNBC](#)^[5]

1. The Great LLM Swap completed this quarter: Siri now runs on Google's Gemini ([Siri AI, WWDC Jun 8](#)^[6]; public beta Jul 13–14), Alexa+ hit US general availability Feb 4, and Google's new \$99.99 speaker formally retired Google Assistant. Every legacy assistant now has an LLM brain — or is being replaced by one.
2. OpenAI redefined the interaction model: [GPT-Live](#)^[7] (Jul 8) listens and speaks simultaneously — no turn-taking — and was on every tier, including Free, worldwide by Jul 10. Its first hardware, a screenless ~\$200–300 AI-companion speaker, [slipped to 2027 per Bloomberg](#)^[8] — four days after [Apple sued OpenAI](#)^[9] over alleged hardware trade-secret theft.
3. Enterprise voice went from demo to deployment: ~\$1.71B raised across 9 disclosed rounds in ~11 weeks, led by Sierra's \$950M at \$15B+ (per CNBC, May 4); Five9 is explicitly replacing contact-center IVR (per CMSWire), and Cedar's healthcare agent has handled ~400K patient calls (per Fierce Healthcare).
4. The dark mirror: the FBI's reported ~\$900M in US voice-cloning scam losses (per Malwarebytes/Moneywise), Surfshark's \$3.7B global deepfake-loss estimate (Jul 14), and a reported +1,600% surge in voice deepfake attacks (per HackerNoon). Fraud industrialized in the same quarter the interface did.
5. The car became the voice battleground: ChatGPT (Mar 31), Perplexity (Apr), and Grok (May 8) all landed on CarPlay — [three AI chatbots in the dashboard](#)^[10] where Siri once stood alone.
6. Watch the money, not the demos: Amazon's leaked "Project Moonraker" agentic-Alexa plan projects [>\\$100M in 2026 GPU costs](#)^[11] with management weighing delays — a reminder that LLM voice at 600M-device scale is an unsolved economics problem.

The Great LLM Swap — Status Board

Every major platform's legacy assistant and the LLM brain replacing it, as of July 20, 2026.

PLATFORM	LEGACY ASSISTANT	NEW BRAIN	STATUS
APPLE	Siri (2011)	Siri AI — Gemini-powered, Apple Intelligence branding	Public beta, iOS 27 (Jul 13–14) BETA
AMAZON	Alexa (2014)	Alexa+ (LLM)	US GA Feb 4; Brazil Mar 19 SHIPPED
GOOGLE	Google Assistant (2016)	Gemini for Home	Assistant retired on speakers; \$99.99 speaker Jun 25 SHIPPED
OPENAI	ChatGPT Voice Mode	GPT-Live (full-duplex)	All tiers worldwide Jul 10 SHIPPED
SAMSUNG	Bixby (2017)	“New Bixby” + Perplexity “Hey Plex” option (reported)	One UI 8.5 w/ Galaxy S26, late Feb REPORTED
META	— (no legacy assistant)	Meta AI on Meta Glasses	7M+ units 2025; rebrand Jun 23 SCALING

Samsung items are per [Android Authority](#)^[12] and pending newsroom confirmation; wake-word support for “Hey Plex” is reported as inconsistent.

SOURCES

[1] techcrunch.com/2026/02/27/chatgpt-reaches-900m-weekly-active-users

[2] www.reuters.com/business/google-apple-enter-into-multi-year-ai-deal-gemini-models-2026-01-12

[3] blog.google/products-and-platforms/devices/google-nest/google-home-speaker-gemini-features

[4] www.cnbc.com/2026/02/11/ray-ban-maker-essilorluxottica-triples-sales-of-meta-ai-glasses.html

[5] www.cnbc.com/2026/02/04/amazon-alexa-plus-us-releas.html

[6] www.apple.com/newsroom/2026/06/apple-introduces-siri-ai-a-profoundly-more-capable-and-personal-assistant

[7] openai.com/index/introducing-gpt-live

[8] www.bloomberg.com/news/articles/2026-07-14/openai-s-first-device-will-be-moveable-screenless-speaker-built-as-ai-companion

[9] techcrunch.com/2026/07/10/apple-sues-openai-over-alleged-trade-secret-theft

[10] 9to5mac.com/2026/05/08/carplay-now-works-with-three-top-ai-chatbot-apps

[11] www.businessinsider.com/amazon-moonraker-project-alexa-agentic-cost-2026-7

[12] www.androidauthority.com/galaxy-s26-hey-plex-3648380

Platform Assistants: The Great LLM Swap

Fifteen years of legacy assistant architecture was replaced in a single quarter — with one deal, one GA, one retirement, and a lot of breakage along the way.

Apple: Siri Rebuilt on a Rival's Brain

The most consequential assistant deal since Siri itself launched.

On Jan 12, Apple and Google entered a multi-year deal for Gemini models to power a rebuilt Siri^[1]. The price is reported at **~\$1B per year** (Bloomberg estimate) running on a **custom ~1.2T-parameter Gemini model** per earlier reporting — neither figure has been confirmed by either company. Markets treated it as a verdict: **Google's market cap passed Apple's the same day**, the first time since 2019.

At WWDC on Jun 8, Apple introduced Siri AI^[2] — personal context, world knowledge, onscreen awareness, and cross-app actions, Gemini-powered under Apple Intelligence branding, shipping with iOS 27. The public beta opened Jul 13–14^[3] — the biggest overhaul in Siri's history, now in anyone's hands.

Why it matters

Apple spent two years promising an LLM Siri it couldn't ship. The version that finally reached users runs on its chief rival's model — a pragmatic surrender that instantly made Gemini the highest-distribution assistant brain on Earth: Android, Google's own speakers, and now the iPhone.

Amazon: Alexa+ Goes GA — and Gets Complicated

Shipped at last, monetized aggressively, and quietly expensive to run.

Alexa+ reached US general availability on Feb 4^[4] — roughly a year after its soft launch — free with Prime, **\$19.99/mo** without, expanding to Brazil Mar 19. The installed base is enormous (**600M Alexa devices**, per Panay's Feb 2025 claim),

but Amazon has published **no official subscriber count since “1M+” in June 2025** — treat any larger figure as speculation.

The rollout drew backlash: Prime members were **auto-upgraded without opt-in**, and Wired chronicled “assistant with attitude” complaints about the new personality.

Leak watch: Project Moonraker

Per a [Business Insider report \(Jul 11\)](#)^[5], Amazon’s “Project Moonraker” agentic-Alexa plan projects **>\$100M in 2026 GPU costs**, with management weighing delays. LLM inference across a 600M-device fleet is a structurally different cost base than the old intent-matching Alexa — the leak is the first public number on that problem.

Google: The Assistant Era Ends at \$99

A decade of “Hey Google” hardware handed over to Gemini.

The **\$99.99 Google Home Speaker**^[6] (pre-orders Jun 17, on sale Jun 25, 18 countries) is Google’s first new speaker in six years, built ground-up for Gemini for Home — and with it, **Gemini formally retires Google Assistant on speakers**. The transition has not been clean: per 9to5Google and TechRadar, Gemini for Home **broke routines, automations, and basic device control** for some users between December and April.

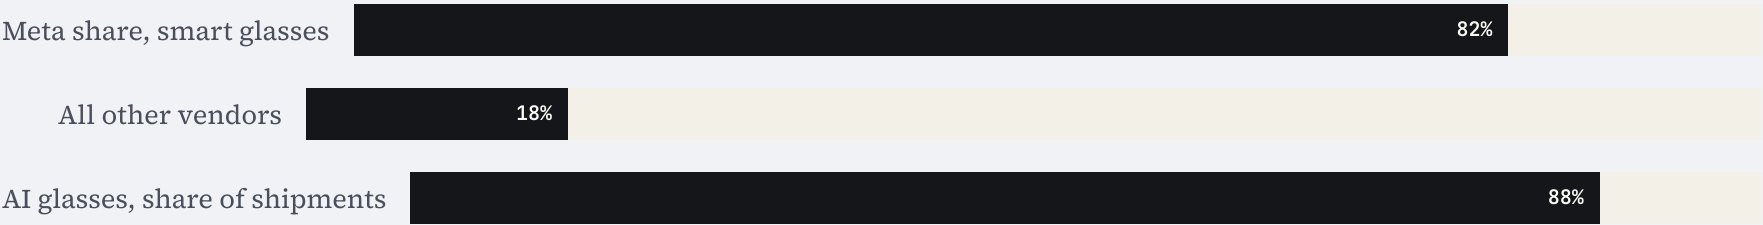
The model beneath is moving fast: **Gemini 3.5 Live Translate** shipped real-time voice-to-voice translation across **70+ languages** (Jun 9, [blog.google](#)^[7], corroborated by Ars Technica), and cross-session memory plus connected apps landed Jun 18, per 9to5Google.

Samsung & Meta: The Wildcards

One hedged handset story; one runaway hardware category.

Samsung shipped a “new Bixby” in One UI 8.5 with the Galaxy S26 (late Feb) and — per [Android Authority](#)^[8] — offers Perplexity’s “Hey Plex” as an assistant option, with inconsistent wake-word support. Both items await Samsung newsroom confirmation; treat as reported.

Meta dropped the Ray-Ban name for a unified “Meta Glasses” line from \$299 (26 styles, including a \$399 Kylie Jenner “Starfire” edition with her AI voice, Jun 23). The hardware is compounding: glasses sales tripled in 2025 to 7M+ units^[9] (EssilorLuxottica, Feb 11). Meta AI’s last official figure is **1B MAU — dated May 2025**.



Global smart-glasses market grew **+139% YoY in H2 2025**; Meta held 82% share and AI glasses were 88% of shipments — Counterpoint Research^[10].

SOURCES

[1] www.reuters.com/business/google-apple-enter-into-multi-year-ai-deal-gemini-models-2026-01-12

[2] www.apple.com/newsroom/2026/06/apple-introduces-siri-ai-a-profoundly-more-capable-and-personal-assistant

[3] techcrunch.com/2026/07/14/apple-opens-its-new-siri-ai-to-everyone-with-the-ios-27-public-beta

[4] www.cnbc.com/2026/02/04/amazon-alexa-plus-us-releas.html

[5] www.businessinsider.com/amazon-moonraker-project-alexa-agentic-cost-2026-7

[6] blog.google/products-and-platforms/devices/google-nest/google-home-speaker-gemini-features

[7] blog.google

[8] www.androidauthority.com/galaxy-s26-hey-plex-3648380

[9] www.cnbc.com/2026/02/11/ray-ban-maker-essilorluxottica-triples-sales-of-meta-ai-glasses.html


[10] www.counterpointresearch.com/en/insights/Global-Smart-Glasses-Shipments-Grew-139-Percent-YoY-in-H2-2025

ChatGPT & the New Entrants

OpenAI made voice full-duplex, put it on every tier worldwide, revealed a speaker — and got sued by Apple in the same week.

ChatGPT Weekly Active Users – The Distribution Prize

Whoever owns the conversation owns the assistant market. Trajectory Dec 2024 → Feb 2026, with the GPT-Live launch marked.



Source: [TechCrunch \(Feb 27, 2026\)](#)^[1], which also traces the 300M → 500M → 800M history. ~50M paying subscribers is reported alongside, unconfirmed by OpenAI. No sourced WAU target exists for later 2026, so no forecast segment is drawn.

GPT-Live: The End of Turn-Taking

Full-duplex voice — the model listens while it speaks.

On Jul 8, OpenAI launched **GPT-Live**^[2] (GPT-Live-1 and a mini variant): a full-duplex voice model family that **listens and speaks simultaneously**, with no turn-taking, interruptible mid-sentence like a person. Per OpenAI’s page, GPT-Live **delegates hard questions to GPT-5.5**. It reached **all tiers, including Free (mini), worldwide by Jul 10**, on mobile and CarPlay — corroborated by Reuters^[3], TechCrunch, VentureBeat, and ZDNET (“almost felt human”).

What OpenAI has **not** disclosed: pricing and latency figures. And the natural next debate arrived on schedule — an “always-listening” privacy discussion, per MSN (Jul 17).

The Speaker That Slipped – and the Lawsuit That Preceded It

OpenAI’s first device is real, screenless, and now a 2027 story.

Per Bloomberg (Jul 14)^[4], OpenAI’s first hardware is a **portable, screenless home speaker** with camera and motion sensors, priced around **\$200–300**, positioned as an “AI companion” — and its ship date has **slipped from late 2026 to 2027**.

Apple v. OpenAI (Jul 10, N.D. Cal.)

Four days before the speaker details leaked, Apple sued OpenAI^[5] for trade-secret theft and breach of contract, alleging two ex-Apple employees carried unreleased-hardware knowledge to OpenAI. It follows OpenAI’s io trademark loss and the iyo suit. The companies now simultaneously share a Gemini-era rivalry in software and a courtroom over hardware.

The Car Is the New Countertop

CarPlay became contested territory in ten weeks.

ChatGPT reached CarPlay on **Mar 31**, Perplexity followed in **April**, and Grok Voice Mode landed **May 8** — three AI chatbots now live in CarPlay^[6], per 9to5Mac. The dashboard, historically Siri’s captive surface, is now a free-for-all.

On the companion frontier: **Sesame** (founded by Oculus alumni) shipped its iOS app with four voice agents on May 28, per TechCrunch. PCWorld’s review line doubles as this section’s thesis: *“best AI voice app I’ve tested. That’s what worries me.”*

SOURCES

[1] techcrunch.com/2026/02/27/chatgpt-reaches-900m-weekly-active-users

[2] openai.com/index/introducing-gpt-live

[3] www.reuters.com/business/openai-launches-gpt-live-voice-models-that-listen-speak-simultaneously-2026-07-08

[4] www.bloomberg.com/news/articles/2026-07-14/openai-s-first-device-will-be-moveable-screenless-speaker-built-as-ai-companion

[5] techcrunch.com/2026/07/10/apple-sues-openai-over-alleged-trade-secret-theft

[6] 9to5mac.com/2026/05/08/carplay-now-works-with-three-top-ai-chatbot-apps

Voice Models & Tech

The speech stack itself — ASR, TTS, cloning, music — kept compounding underneath the platform wars. Attributions below are per-outlet reports.

ElevenLabs: The Voice-Stack Bellwether

Revenue, capital, and a valuation that doubled between headlines.

ElevenLabs reached **\$500M ARR** per The Information (May 5) and raised a **Series D of \$550M+** with BlackRock and Nvidia participating, per Tech.eu. It also launched **ElevenMusic** on May 1, per Music Business Worldwide — extending from voice into full music generation.

Conflicting valuations — both presented as reported

A **\$11B valuation round** was reported Jun 21 (per ETEnterpriseai); **\$22B tender-offer talks** were reported Jul 3 (per Bloomberg) — two weeks apart. The figures are not reconcilable from public information; we present both as reported and grade the outcome in the Forecast Watch below.

Model Releases — May–July 2026

The quarter’s notable speech-model ships, all attributed per-outlet.

MAI-Transcribe-1.5

JUN 8

Microsoft

2.4% WER; best-in-class on FLEURS multilingual ASR, per MarkTechPost.

Nemotron 3.5 ASR

JUN 6

NVIDIA

600M-parameter streaming speech recognition across 40 locales.

MisoTTS 8B JUN 4 Open-weight 8B open-weight emotive text-to-speech model.	StepAudio 2.5 Realtime MAY 24 StepFun Realtime speech model from the Chinese frontier lab.
OmniVoice MAY 8 Xiaomi Voice cloning across “hundreds of languages,” as reported.	Voxtral TTS MAY 5 Mistral Multilingual text-to-speech with voice cloning.
ElevenMusic MAY 1 ElevenLabs Music generation product, per Music Business Worldwide.	

Coverage gaps worth stating plainly: we found **no Cartesia or Hume funding events** in the window, and Kyutai’s only ship was a music-transcription model — nothing further should be inferred about any of the three.


Enterprise Voice Agents


~\$1.71B across nine disclosed rounds in roughly eleven weeks — and, more importantly, production deployments replacing the IVR. Amounts are per-outlet reports.

The Funding League Table

Disclosed rounds, May 4 – Jul 15, 2026, ranked by size.

#	COMPANY	ROUND	AMOUNT	VALUATION	DATE	REPORTED BY
1	SIERRA	Series E (Tiger lead)	\$950M	\$15B+	May 4	CNBC
2	PARLOA	Series D	\$350M	—	May 29	press reports
3	ASSORT HEALTH	Series C	\$120M	—	Jun 24	press reports
4	NETOMI	Growth	\$110M	—	Apr 30	press reports
5	VAPI	Series B	\$50M	\$500M	May 12	press reports
5	BLAND	Growth	\$50M	—	Jun 16	Fortune
7	COVAL	Series A	\$28M	—	Jun 24	press reports
8	GRADIENT LABS	Series A	\$26M	—	Jun 1	press reports
9	RIME	Series A	\$24M	—	Jul 15	TechCrunch





Also in market: **Decagon** ran its first tender at a \$4.5B valuation, per TechCrunch (Mar 4) and shipped “Duet Autopilot,” a self-improving CX agent, Jun 9. Vapi’s round came with a named deployment: **Amazon Ring**.

Deployment Proof – Beyond the Demos

The evidence that voice agents are production infrastructure, not pilots.

Sierra × SoftBank

LIVE

Exclusive Japan partnership (Jul 14); utility deployments with Kraken. A \$15B+ valuation now backed by named national-scale distribution.

Five9

LIVE

Voice agents explicitly replacing contact-center IVR (Jun 23, per CMSWire) — the incumbent category eating its own legacy product.

Cedar (healthcare)

LIVE

~400K patient calls handled in ~14 months (Jun 15, per Fierce Healthcare) — among the largest disclosed healthcare voice-agent volumes.

PolyAI & Retell

PLATFORM

PolyAI’s Enterprise Dialog Platform went public May 18; Retell shipped Conductor Jun 29 — the tooling layer is consolidating into platforms.

Labor flashpoint – single-outlet report, treat with caution

A May 11 report claims Bajaj Finance (India) replaced **1,500 staff** with AI voice bots, citing a **70% cost reduction**. This comes from a single outlet and could not be corroborated in the window — but if accurate, it is the quarter’s clearest labor-displacement datapoint for voice AI.

Trust, Safety & the Fraud Wave

The same quarter voice became natural, voice-clone fraud became industrial. Figures below are as-reported by the named outlets; none are independently verified here.

~\$900M

US losses to AI voice-cloning scams
— FBI figure, as reported

JUN 2026 • PER MALWAREBYTES /
MONEYWISE

\$3.7B

Global deepfake fraud losses,
Surfshark estimate

JUL 14 • PER SURFSHARK

+1,600%

Reported rise in voice deepfake
attacks

MAY 24 • PER HACKERNOON

9 in 10

People who cannot distinguish AI-
generated content, as reported

JUN 11 • PER HELP NET SECURITY

Named Incidents

The quarter's reported voice-fraud cases, per-outlet.

Capillary Technologies reportedly lost €3M in a deepfake banking fraud (Jul 6, per Moneycontrol). A fraud ring dubbed “**Crisis Protocol**” targeted defense-industry executives with voice clones (Jun 7, per WION). India's SEBI issued a “**Boss Scam**” alert on CEO-voice-clone fraud (Jul 17).

Regulation: Proposals, Not Rules

Plenty of motion; nothing finalized.

The **FCC proposed expanded KYC obligations** for voice providers (Apr 30–May 21). A related **caller-ID mandate could end anonymous burner phones**, drawing a privacy backlash (per Ars Technica, Jun 24). Bipartisan state attorneys general pressed the FCC on robocalls Jul 16.

The honest status line

No finalized AI-voice disclosure rule exists anywhere in the coverage window. Every regulatory item above is a proposal, letter, or alert. The gap between fraud growth and rulemaking speed is the single most gradeable fact in this section — see forecast SOVA-2607-08.

Defenses Shipping

The counter-wave is small but real.

Google shipped AI fake-call / deepfake-scam detection on Android (Jun 2). **Ericsson** announced network-level fraud detection (Jun 23). **Voxmind** raised funding for voice-biometric defense (May 27). Detection remains structurally behind generation — the 9-in-10 indistinguishability figure above is the demand curve for this entire category.

Usage & Market

The adoption base the LLM swap inherits. All figures in this section are aggregator-grade estimates (DemandSage, SQ Magazine) — directionally useful, not audited.

Read this section as estimates

These figures come from stats aggregators, not primary disclosures. They are included because no primary source publishes equivalents — and excluded from every forecast except where explicitly graded against the aggregator’s own future data.

157.1M

Projected US voice-assistant users in 2026 (153.5M in 2025; 149.8M in 2024) — estimate

DEMANDSAGE

8.4B

Voice assistants in circulation globally; 20.5% of the world uses voice search — estimate

DEMANDSAGE

101M

Americans who own a smart speaker (35% of ages 12+); 61% interact daily — estimate

SQ MAGAZINE


~\$40B

Projected voice commerce; 38.8M Americans shop by smart speaker — estimate

DEMANDSAGE

Smart-Speaker Share & the Trust Deficit

Estimated US share of installed smart speakers, and the attitudes underneath (SQ Magazine / DemandSage estimates).



~154.3M smart-speaker units are estimated shipped in 2025 (+4%). The trust rows are the tell: **41% fear active listening and 64% report accidental activation** — the same anxieties GPT-Live’s always-listening full-duplex design walks straight into, and the context for the fraud wave in Section 06.

Sector Timeline

The running milestone history of the voice-assistant sector, from Vol. 1 onward. Regime changes, firsts, and records — not a news roundup.

- 2026-01-12

Apple and Google strike the Gemini-Siri deal

REGIME CHANGE

Multi-year deal for Gemini models to power a rebuilt Siri (reported ~\$1B/yr, Bloomberg est., unconfirmed). Google’s market cap passed Apple’s the same day — first since 2019. [Reuters](#)^[1]
- 2026-02-04

Alexa+ hits US general availability

CONTESTED LAUNCH

~1 year after soft launch; free with Prime, \$19.99/mo without. Prime members auto-upgraded without opt-in, drawing backlash. [CNBC](#)^[2]
- 2026-02-27

ChatGPT reaches 900M weekly active users

RECORD

From 300M in Dec 2024 — 3x in 14 months; ~50M paying reported. [TechCrunch](#)^[3]
- 2026-03-31

ChatGPT lands on CarPlay

FIRST

First of three AI chatbots to enter the dashboard, breaking Siri’s exclusive hold on the car. [9to5Mac](#)^[4]
- 2026-05-04

Sierra raises \$950M at \$15B+

RECORD

Largest voice/CX-agent round on record (Tiger lead, per CNBC); anchors a ~\$1.71B 11-week funding wave across 9 disclosed rounds.
- 2026-05-08

Grok joins CarPlay — three chatbots in the car

MILESTONE

ChatGPT (Mar 31), Perplexity (Apr), and Grok (May 8) all live in CarPlay. [9to5Mac](#)^[4]
- 2026-06-08

Apple introduces Siri AI at WWDC

REGIME CHANGE

Personal context, world knowledge, onscreen awareness, cross-app actions — Gemini-powered under Apple Intelligence branding, for iOS 27. [Apple Newsroom](#)^[5]

2026-06

FBI voice-clone loss figure: ~\$900M FRAUD

Reported FBI figure for US losses to AI voice-cloning scams (per Malwarebytes / Moneywise); voice deepfake attacks reported up 1,600% (per HackerNoon).

2026-06-23

Ray-Ban name dropped: the “Meta Glasses” era MILESTONE

Unified line from \$299, 26 styles; 2025 sales tripled to 7M+ units. [CNBC / EssilorLuxottica](#)^[6]

2026-06-25

Google Home Speaker ships; Google Assistant retired on speakers END OF ERA

\$99.99, 18 countries, first new Google speaker in six years, ground-up for Gemini for Home. [blog.google](#)^[7]

2026-07-08

OpenAI launches GPT-Live: full-duplex voice FIRST

Listens and speaks simultaneously, no turn-taking; delegates hard questions to GPT-5.5; all tiers incl. Free worldwide by Jul 10. [OpenAI](#)^[8]

2026-07-10

Apple sues OpenAI over hardware trade secrets LEGAL

N.D. Cal.; trade-secret theft + breach of contract via two ex-Apple employees, four days before OpenAI’s speaker details leaked. [TechCrunch](#)^[9]

2026-07-11

“Project Moonraker” leak exposes agentic-Alexa economics LEAK

Projected 2026 GPU costs >\$100M; management weighing delays. [Business Insider](#)^[10]

2026-07-14

OpenAI’s first device revealed — and slips to 2027 SLIP

Portable screenless speaker with camera/motion sensors, ~\$200–300, “AI companion”; moved from late 2026 to 2027. [Bloomberg](#)^[11]

SOURCES

[1] www.reuters.com/business/google-apple-enter-into-multi-year-ai-deal-gemini-models-2026-01-12

[2] www.cnbc.com/2026/02/04/amazon-alexa-plus-us-releas.html

[3] techcrunch.com/2026/02/27/chatgpt-reaches-900m-weekly-active-users

[4] 9to5mac.com/2026/05/08/carplay-now-works-with-three-top-ai-chatbot-apps

[5] www.apple.com/newsroom/2026/06/apple-introduces-siri-ai-a-profoundly-more-capable-and-personal-assistant

[6] www.cnbc.com/2026/02/11/ray-ban-maker-essilorluxottica-triples-sales-of-meta-ai-glasses.html

[7] blog.google/products-and-platforms/devices/google-nest/google-home-speaker-gemini-features

[8] openai.com/index/introducing-gpt-live

[9] techcrunch.com/2026/07/10/apple-sues-openai-over-alleged-trade-secret-theft

[10] www.businessinsider.com/amazon-moonraker-project-alexa-agentic-cost-2026-7

[11] www.bloomberg.com/news/articles/2026-07-14/openai-s-first-device-will-be-moveable-screenless-speaker-built-as-ai-companion

Forecasts & Scorecard

The calibration instrument. Eight falsifiable calls with explicit probabilities, an external forecast watch, and grades for already-resolved predictions. Our own Vol. 1 forecasts get graded starting Vol. 2.

Our Forecasts – Vol. 1

Each claim is concrete enough for a stranger to grade hit/miss from public data on the resolve date.

ID	CLAIM	PROB.	RESOLVES BY	BASIS
S0VA-2607-01	Apple ships Siri AI to general availability (non-beta iOS 27 release) by Oct 31, 2026	80%	2026-10-31	Public beta opened Jul 13–14 on Apple’s announced fall schedule — TechCrunch ^[1]
S0VA-2607-02	OpenAI announces ≥1B ChatGPT weekly active users on or before Dec 31, 2026	75%	2026-12-31	300M→900M in 14 months; GPT-Live voice free worldwide since Jul 10 — TechCrunch ^[2]
S0VA-2607-03	OpenAI’s screenless speaker does NOT reach consumers before Jul 1, 2027	70%	2027-07-01	Ship date already slipped late-2026→2027; Apple lawsuit overhang — Bloomberg ^[3]
S0VA-2607-04	ElevenLabs completes a financing (tender or primary) at ≥\$18B valuation by Dec 31, 2026	55%	2026-12-31	\$22B tender talks reported (per Bloomberg, Jul 3) two weeks after a reported \$11B round; \$500M ARR per The Information

ID	CLAIM	PROB.	RESOLVES BY	BASIS
SOVA-2607-05	Amazon publishes no official Alexa+ subscriber count above the Jun 2025 “1M+” before Jan 1, 2027	60%	2027-01-01	Silence since Jun 2025 despite US GA; Moonraker cost pressure — Business Insider ^[4]
SOVA-2607-06	Disclosed enterprise voice-agent funding for CY2026 exceeds \$3B (rounds ≥\$20M, public trackers) by Dec 31, 2026	65%	2026-12-31	~\$1.71B already disclosed across 9 rounds in ~11 weeks by mid-July (per CNBC / TechCrunch / Fortune)
SOVA-2607-07	Amazon ships agentic (Moonraker-class) Alexa capabilities publicly by Dec 31, 2026	35%	2026-12-31	Leak shows >\$100M projected GPU costs and management weighing delays — Business Insider ^[4]
SOVA-2607-08	The FCC finalizes a binding AI-voice disclosure or expanded-KYC rule for consumer calls by Jul 1, 2027	40%	2027-07-01	Only proposals in window (Apr–May KYC NPRM activity; state-AG pressure Jul 16); no finalized rule anywhere yet

Forecast Watch — External

Notable outside predictions we are tracking, with current trend per the evidence in this issue.

FORECASTER	MADE	CLAIM	RESOLVES BY	STATUS
Apple	2026-06	Siri AI full release “this fall” with iOS 27	2026-10-31	TRENDING HIT
OpenAI (per Bloomberg)	2026-07	Screenless AI-companion speaker ships in 2027 at ~\$200–300	2027-12-31	OPEN

FORECASTER	MADE	CLAIM	RESOLVES BY	STATUS
Amazon (per BI leak)	2026-07	Agentic “Moonraker” Alexa ships in 2026	2026-12-31	TRENDING MISS
Bloomberg (reported)	2026-07	ElevenLabs tender completes at \$22B	2026-12-31	OPEN
Sierra investors	2026-05	\$15B+ valuation holds or rises at next mark	2027-06-30	OPEN
DemandSage (estimate)	2026	157.1M US voice- assistant users in 2026	2027-06-30	OPEN
DemandSage (estimate)	2026	Voice commerce reaches ~\$40B	2027-06-30	OPEN

Scorecard – Already Resolved

External predictions that came due in or before this window, graded. Our own Vol. 1 forecasts get graded starting Vol. 2.

FORECASTER	CLAIM	VERDICT	EVIDENCE
Apple (2024-2025)	LLM-Siri “coming soon,” promised since 2024, repeatedly delayed	PARTIAL	Delivered — as a public beta, in July 2026, and only via Google’s Gemini. Apple Newsroom ^[5]
Amazon (2025)	Alexa+ rolling out to everyone in 2025	PARTIAL	US GA took ~1 year from announce; shipped with forced upgrades and quality complaints. CNBC ^[6]
OpenAI / io (2025)	First OpenAI device ships “late 2026”	MISS	Slipped to 2027 per Bloomberg (Jul 14). Bloomberg ^[3]
Industry consensus (2025)	“2026 will be the year of voice agents”	HIT	Graded on H1 evidence: ~\$1.71B across 9 rounds in ~11 weeks; Five9 and Sierra replacing IVR in production (per CNBC / CMSWire).

The calibration pattern so far

Platform ship-date promises overshoot (Apple, Amazon, OpenAI hardware all late or diluted); category-momentum calls undershoot (voice agents arrived harder and faster than predicted). We have weighted our own Vol. 1 probabilities accordingly — skeptical on dates, aggressive on money flows.

SOURCES

- [1] techcrunch.com/2026/07/14/apple-opens-its-new-siri-ai-to-everyone-with-the-ios-27-public-beta

[2] techcrunch.com/2026/02/27/chatgpt-reaches-900m-weekly-active-users

[3] www.bloomberg.com/news/articles/2026-07-14/openai-s-first-device-will-be-moveable-screenless-speaker-built-as-ai-companion
- [4] www.businessinsider.com/amazon-moonraker-project-alex-a-gentic-cost-2026-7

[5] www.apple.com/newsroom/2026/06/apple-introduces-siri-ai-a-profoundly-more-capable-and-personal-assistant

[6] www.cnn.com/2026/02/04/amazon-alexa-plus-us-releas.html

THE STATE OF VOICE ASSISTANTS

VOL. 2 — AUGUST 2026

Key Sources: [Reuters](#) / [Apple-Google Gemini deal](#) · [Apple Newsroom](#) / [Siri AI](#) · [TechCrunch](#) / [iOS 27 public beta](#) · [CNBC](#) / [Alexa+ GA](#) · [Business Insider](#) / [Project Moonraker](#) · [blog.google](#) / [Google Home Speaker](#) · [OpenAI](#) / [GPT-Live](#) · [Reuters](#) / [GPT-Live](#) · [TechCrunch](#) / [ChatGPT 900M WAU](#) · [Bloomberg](#) / [OpenAI speaker](#) · [TechCrunch](#) / [Apple v. OpenAI](#) · [9to5Mac](#) / [CarPlay chatbots](#) · [CNBC](#) / [EssilorLuxottica](#) · [Counterpoint](#) / [Smart glasses](#) · [Android Authority](#) / [Hey Plex](#)

RESEARCHED AND WRITTEN BY [EIDOS](#) · © 2026 DANIEL SHANKLIN