

AI News

Trends — January 2026

1 Executive Summary

January 2026 was a transitional month. No flagship *frontier model* reached general availability — the period operated principally on the Q4 2025 model fleet (*Claude* Opus 4.5, Sonnet 4.5, Haiku 4.5; **OPENAI** *GPT*-5.2; **GOOGLE** Gemini 3 Pro and Gemini 3 Flash; **MISTRAL** Large 3; **DEEPSEEK** V3.2). Activity concentrated in three domains: product and feature launches, funding deployment, and regulatory effective dates.

ANTHROPIC launched *Claude Cowork* in January — a persistent *agent*-driven workflow product targeting professional use cases such as legal and financial analysis. The company stated during its February 12 Series G announcement that more than thirty products and features had launched at **ANTHROPIC** during January alone. The *Cowork* release coincided with **ANTHROPIC** becoming a **MICROSOFT** subprocessor on January 6, enabling later integration of *Cowork* technology into **MICROSOFT** 365 Copilot. *Claude Code*, which had been generally available since May 2025, went viral over the 2025-2026 winter holidays, expanding from a developer tool into broader "*vibe coding*" use by non-programmers. Weekly active *Claude Code* users doubled between January 1 and February 12.

Funding activity was unusually concentrated. **xAI** closed a \$20 billion Series E led by **NVIDIA**, with participation from Fidelity, the Qatar Investment Authority, and Valor Equity Partners. The round took **xAI**'s valuation past \$230 billion. Independent estimates placed cumulative AI fundraising at approximately \$47 billion in the first two weeks of January alone. **ANTHROPIC**'s Series G negotiation — which would close on February 12 at \$30 billion and a \$380 billion post-money valuation — was actively underway throughout the month.

The *Texas Responsible Artificial Intelligence Governance Act* took effect on January 1, becoming the most consequential US state AI regulation in force at the start of 2026. The European Commission continued negotiation on the *Digital Omnibus on AI* — adopted in November 2025 — in advance of the 2 August 2026 *EU AI Act* enforcement deadline.

To watch in February: **ANTHROPIC**'s expected Opus 4.6 release on roughly the company's two-month cadence, **OPENAI**'s expected *GPT*-5.3 line continuation, **GOOGLE**'s expected Gemini 3.1 release, the closing of the **ANTHROPIC** Series G, and the **OPENAI** mega-round expected to break \$100 billion.

2 Large Language Models

January *LLM* activity centred on incremental updates and platform integration rather than flagship-tier releases. The major frontier models in active deployment had all shipped during Q4 2025; January was their first full month of widespread enterprise rollout.

2.1 Releases and updates

No frontier-tier general-availability model release occurred during January. The largest release activity sat in product and platform integration. **ANTHROPIC** shipped more than thirty products and features during the month,

the most active product release period in the company's history. The principal release was *Claude Cowork*.

OPENAI continued the *GPT-5.2* line throughout January with multiple ChatGPT-side updates and expansion of Codex availability. The *GPT-5.3-Codex* variant would not arrive until February 5; January was the last month of the unified *GPT-5.2* default at the consumer tier. **GOOGLE** maintained Gemini 3 Pro as the flagship and Gemini 3 Flash (released December 17, 2025) as the cost-efficient variant. **MISTRAL** maintained **MISTRAL** Large 3, released in December 2025, alongside the Ministral 3 family.

DEEPSEEK V3.2-Exp, released in late 2025 as a sparse-attention experimental architecture, continued to operate as the principal cost-disruption *open-weight* option — reaching frontier-comparable *benchmark* performance at a fraction of the per-*token* cost of closed proprietary models. The **DEEPSEEK** pricing structure, established through V3.2, reset the cost floor across the segment heading into 2026.

2.2 Benchmarks and capabilities

Independent *benchmark* trackers entering January reflected a competitive landscape in which *Claude* Opus 4.5 led on *SWE-bench* Verified at 80.9%, *GPT-5.2* led on certain reasoning subsets, and Gemini 3 Pro led on *multimodal* evaluations. The *Artificial Analysis Intelligence Index*, with the v4.0 methodology fully in force, measured Gemini 3 Pro Preview at 56 points and *GPT-5.2* (xhigh) at 51 points.

Long task-completion horizons emerged as the most discussed differentiator at the start of the year. *METR*'s task-horizon estimates placed *Claude* Opus 4.5 at approximately 8 hours at the 50% mark — the highest among the November 2025 cohort. The metric reflected the practical consideration that frontier-class *agentic* deployment requires multi-hour coherence, not single-step accuracy.

Computer use remained below the human expert baseline at the start of January. Both *Claude* Sonnet 4.5 and **OPENAI** Codex variants approached but did not exceed the 72.4% *OSWorld-Verified* threshold. The first model to cross that threshold (*GPT-5.4* at 75% on March 5) was still two months away.

2.3 Pricing and access

Pricing remained stable through January at the levels established with the Q4 2025 cohort: *Claude* Opus 4.5 at \$5 per million input and \$25 per million output *tokens*; Sonnet 4.5 at \$3 / \$15; Gemini 3 Pro and Flash at the existing tier; *GPT-5.2* across multiple sub-variants. The 1-million-*token context window* remained a beta-only feature on *Claude* Sonnet 4.5 (with the *long-context* premium pricing for requests over 200K *tokens*), and was not yet on Opus 4.5.

Free-tier access continued primarily through ChatGPT (limited *GPT-5.2* access), *Claude* (Sonnet via the free tier), and Gemini (Flash variants). **DEEPSEEK** V3.2-Exp through *API* providers and direct hosting remained the principal cost-disruption alternative.

3 Generative Media

Generative media activity in January concentrated in iterative product updates rather than flagship architecture releases. The principal Q1 video releases (Kling 3.0, Seedance 2.0, Sora 2 Pro) would arrive in February.

3.1 Image

No major flagship image model release occurred in January. The active flagship cohort — **BLACK FOREST LABS** FLUX 1.1, **STABILITY AI** SD3.5, **ADOBE** Firefly Image, Midjourney V7, **GOOGLE** Imagen 4 — continued through iterative updates and integration into editor workflows. The pattern observed across late 2025 of capability moving inside editor environments rather than through standalone model releases continued.

3.2 Video

LTX 2.0 was released in January as an open-source *text-to-video* model supporting resolutions from 1080p to 4K. The release continued the open-source pressure on proprietary video generation and represented one of the few flagship-tier releases of the month in any segment.

RUNWAY shipped a Gen-4.5 update during January, focusing on *image-to-video* for longer narrative sequences and **ADOBE** Firefly integration. **GOOGLE** Veo 3.1 Fast launched in January, providing two-times-faster generation for rapid iteration use cases. The active video cohort — **OPENAI** Sora 2, **RUNWAY** Gen-4.5, **GOOGLE** Veo 3.1, Pika 2.0, Kling 2.x, **BYTEDANCE** Seedance 1.5 — continued through the month.

3.3 Music and audio

SUNO revised its Help Center documentation in January to clarify ownership, commercial use, and copyright responsibility for AI-generated music. The clarifications were not new policy but explicit positioning ahead of distributor and licensing platform questions about authorship. **SUNO**'s approach reflected the maturation of AI music distribution from experimental to commercially registered. Fish Audio, released in January, addressed Asian-language accent retention in voice cloning at production quality. **ELEVENLABS** continued integration partnerships, including the LTX Studio audio-to-video pipeline.

4 Coding and Developer Tools

January was the breakout month for AI-assisted coding. *Claude Code* revenue, **CORSOR** adoption, and **GITHUB** Copilot subscription growth all reached inflection

points coinciding with the holiday-season viral wave that originated in late December 2025.

4.1 Releases

ANTHROPIC launched *Claude Cowork* in January — the company's principal product release of the month. *Cowork* brought *Claude Code*'s *agentic* engineering capabilities to a broader scope of knowledge work tasks, with persistent *agent* threads available to Pro and Max plan users on PC through the *Claude* Desktop app. The initial release targeted professional use cases (legal and financial analysis) with plugin support from launch. The *Cowork* architecture would later be licensed to **MICROSOFT** for use inside **MICROSOFT** 365 Copilot, beginning with Copilot *Cowork* in research preview.

OPENAI continued Codex updates throughout January in advance of the *GPT-5.3-Codex* release in February. **CORSOR** maintained its \$20 per month Pro tier, having reached \$2 billion in annualized run-rate by year-end 2025. **GITHUB** Copilot continued at scale, with **MICROSOFT**'s enterprise distribution placing it in the largest installed base of any AI coding product. **REPLIT** continued growth in advance of its March Series D.

4.2 Pricing and licensing

Pricing remained stable through January. Devin's tiered pricing, established in late 2025 at \$20 Core plus \$2.25 per *Agent* Compute Unit, continued. The structural pricing pressure during the month came from the *open-weight* side: **DEEPSEEK** V3.2-Exp at frontier-comparable performance at sub-cent per million *tokens* established a cost floor that proprietary tools could not match for cost-sensitive workloads.

4.3 Adoption signals

Claude Code went viral over the 2025-2026 winter holidays. **ANTHROPIC**'s later disclosure (during the February 12 Series G announcement) showed weekly active *Claude Code* users had doubled between January 1 and February 12 — a doubling in six weeks. Approximately 4% of all **GITHUB** public commits worldwide were authored by *Claude Code* by mid-February, double the percentage from one month prior. The viral wave during the holiday period brought significant non-programmer use — “*vibe coding*” — broadening the addressable market for AI coding products beyond professional developers.

ANTHROPIC *Claude Code* revenue reached \$2.5 billion annualized by mid-February (disclosed in the Series G announcement), more than doubling from the level at the start of January. Business subscriptions to *Claude Code* quadrupled over the same six weeks. Enterprise users represented over half of *Claude Code* revenue.

5 Industry Trends

Q1 2026 venture capital deployment opened at an unprecedented pace. Independent estimates placed cumulative AI fundraising at approximately \$47 billion in the first two weeks of January alone, before **xAI**'s \$20 billion Series E and **ANTHROPIC**'s ongoing Series G negotiations were factored in. The Texas AI law took effect on January 1, beginning a year of increasingly active US state-level enforcement.

5.1 Funding and valuations

xAI closed a \$20 billion Series E in early January, led by **NVIDIA** with participation from Fidelity, the Qatar Investment Authority, Valor Equity Partners, and other strategic and institutional investors. The round took **xAI**'s valuation past \$230 billion, against the \$200 billion level reached after a \$10 billion raise in late 2025 coinciding with Grok 4. Total **xAI** funding now exceeded \$42 billion in reported equity and debt. The company's positioning as a foundational lab — rather than a startup — was confirmed by the speed and scale of the round.

ANTHROPIC's Series G negotiation was actively underway throughout January at terms that would close on February 12 at \$30 billion and a \$380 billion post-money valuation. Reported during the month: Sequoia Capital was leading a portion of the round, despite already backing **OPENAI** and **xAI** — a shift away from traditional portfolio rules against backing competitors. **MICROSOFT** committed up to \$5 billion and **NVIDIA** up to \$10 billion to the **ANTHROPIC** round, both portions of the previously announced November 2025 commitments.

OPENAI continued in active fundraising discussions during January for what would close at \$122 billion at the end of March. Reported negotiations during January discussed valuations near \$830 billion. Other significant Q1 deployments — **WAYMO** \$16 billion, **DATABRICKS** \$7 billion, **POLYMARKET** \$2.6 billion, **SHIELD AI** \$2.3 billion — began rolling out during the quarter.

5.2 Regulation and legal

The *Texas Responsible Artificial Intelligence Governance Act* took effect on January 1, 2026. *TRAIGA* establishes disclosure requirements for AI system deployers in Texas, prohibits certain discriminatory uses of AI in employment and housing decisions, and creates state-level enforcement authority. Texas became the second US state with comprehensive AI legislation in force, following the Colorado *AI Act* (which becomes applicable in June 2026).

The European Commission continued negotiation on the *Digital Omnibus on AI* — adopted in November 2025 — in advance of the 2 August 2026 *EU AI Act* enforcement deadline. The Code of Practice on Marking and Labelling of AI-generated Content remained in first-draft form, with the second draft due in March. Article 50 transparency obligations were on track for the original 2 August 2026 schedule.

The principal *generative AI* copyright cases — New York Times v. **OPENAI** and Getty Images v. **STABILITY AI** — continued through pre-trial motions during January. No decisive rulings were issued during the month.

5.3 Mergers, acquisitions, exits

OPENAI did not complete a publicly disclosed acquisition during January but had begun the Q1 acquisition programme that would result in six completed transactions by mid-April. **ANTHROPIC**'s acquisition of **VERCEPT**, the only publicly disclosed **ANTHROPIC** acquisition during 2026, was completed during Q1.

Z.ai (**ZHIPU AI**) and **MINIMAX** progressed toward Hong Kong Stock Exchange listings during the quarter. Both companies were valued above \$6 billion at listing, completed during Q1. The Hong Kong route emerged as the principal IPO venue for Chinese AI labs absent direct US listings.

5.4 Infrastructure

ANTHROPIC became a **MICROSOFT** subprocessor on January 6, 2026 — a regulatory and procurement milestone enabling subsequent integration of *Cowork* technology into **MICROSOFT** 365 Copilot. The change aligned the data-residency posture of **ANTHROPIC**-powered features with the existing **MICROSOFT** enterprise compliance framework, smoothing later deployment to Fortune-class customers.

Compute infrastructure constraints remained the binding capacity for AI deployment. The frontier labs continued to negotiate multi-year commitments with **NVIDIA**, **AMD**, **AWS**, **MICROSOFT**, and **GOOGLE**. The **ANTHROPIC-COREWEAVE** deal, announced later in 2026, was in negotiation during January. The **GOOGLE**-Broadcom *TPU* partnership was already extending. Reports during the month placed projected US electricity shortfall driven by AI data centre construction at 9 to 18 gigawatts by 2027.

6 Monthly Recap

Chronological summary of the month's reportable events. Sources are linked in the corresponding section above.

Date	Category	Event	Source / Impact
January 1	Regulation	Texas Responsible Artificial Intelligence Governance Act takes effect	Second comprehensive US state AI law in force
Early January	Funding	xAI Series E \$20 billion led by NVIDIA	Valuation past \$230 billion

Date	Category	Event	Source / Impact
January 6	Industry	Anthropic becomes a Microsoft subprocessor	Enables Cowork integration into M365 Copilot
January	Coding	Anthropic launches Claude Cowork	Persistent agent threads, professional workflows
January	Coding	Claude Code goes viral over winter holidays	Weekly active users on track to double in 6 weeks
January	Generative media	LTX 2.0 released as open-source text-to-video	1080p to 4K resolution support
January	Generative media	Runway Gen-4.5 update; Google Veo 3.1 Fast; Fish Audio for Asian-language voice	Iterative updates across video and audio
January	Funding	Cumulative AI fundraising estimated at \$47B in first two weeks	Q1 2026 unprecedented deployment pace
January	Funding	Anthropic Series G negotiation active throughout the month	Closes February 12 at \$30B / \$380B post-money
January	Funding	OpenAI fundraising discussions reportedly target \$830B valuation	Closes at \$122B at end of March
January	LLM	Anthropic ships 30+ products and features	Most active product month in company history

7 Outlook

ANTHROPIC's roughly two-month release cadence, with *Claude* Opus 4.5 having shipped on November 1, 2025, suggests a *Claude* Opus 4.6 release in early February. **ANTHROPIC**'s thirty-product January release wave indicates accelerating product velocity ahead of the model release.

OPENAI's release cadence suggests a successor to *GPT-5.2* in early-to-mid February. The expected naming continuation is *GPT-5.3* with a Codex-specialised variant, consolidating coding capabilities in a dedicated line. The Sora consumer product remains under operation; no shutdown signal had emerged from **OPENAI** by the close of January.

GOOGLE's next Gemini iteration is expected in mid-to-late February. The roughly three-month cadence (Gemini 3 Pro in November 2025, Gemini 3 Flash in December 2025, Gemini 3 Deep Think during the Q1 research preview window) suggests a Gemini 3.1 release with refined reasoning capability. The reasoning techniques from the Gemini 3 Deep Think research line are likely to be made commercially accessible.

ANTHROPIC's Series G is expected to close in February at terms above \$350 billion. **OPENAI**'s mega-round, in active fundraising during January, is expected to break \$100 billion. Total Q1 venture deployment will set new quarterly records on independent estimates.

The *Digital Omnibus on AI* is expected to reach formal political agreement during March. The implementation timing of any postponement of high-risk obligations under the *EU AI Act* remains contingent on Council and Parliament negotiation. Companies operating in or selling into the European market should treat the 2 August 2026 deadline as binding.

The principal generative-AI copyright cases continue toward decisive phases. Initial summary judgement rulings remain on track for Q2 2026. No major rulings are expected during February.