



MONTHLY AI-TOOL CHANGES REPORT

20 Verified AI-Tool Changes That Mattered in September 2026

Evidence-based analysis of the launches, capability and access changes, plans, governance and retirement that materially affected AI-tool users during September.

VERSION 1.0

20	14	2	7
material changes verified	model and product launches	capability changes	vendor ecosystems

September 1 to September 30, 2026 | Verified through October 1, 2026 | Chooseely Editorial

Model or product launch		14
Capability change		2
Availability or access		1
Plan change		1
Policy change		1
Shutdown		1

Key findings

September's 20 verified events were dominated by model and product launches. The count describes this selected, decision-relevant dataset, not the whole AI market or every catalog update.

The more consequential shift was in work surfaces. Claude folded Cowork into chat and introduced Docs and Slides; ChatGPT added persistent dots; the Agents API gained hosted browser actions. Buyers now need to compare permission boundaries and accepted work, not a model score alone.

Price signals moved in different ways. GPT-6 Sol and Luna launched with lower API rates, Opus 5.5 cut standard and cache-read rates, and Pro 500 introduced a new US\$500 monthly access tier. These are secondary pricing effects of launches except for the separate plan change.

Availability still qualified every headline: Mythos 5.1 and Gemini Flash Cyber used restricted access, Siri AI had a gradual beta rollout, dots rolled out by plan and market, and enterprise administrators controlled several features.

September at a glance

Primary change class	Changes	Share
Model or product launch	14	70%
Availability or access change	1	5%
Capability change	2	10%
Plan change	1	5%
Price change	0	0%
Policy change	1	5%
Deprecation	0	0%
Shutdown	1	5%

One primary class per event. Secondary consequences do not add headline counts.

Coverage statement

Twenty selected events from OpenAI, Anthropic, Google, GitHub, Apple, Suno and SpaceXAI passed the materiality and first-party evidence threshold. The set is concentrated and is not a census of the AI market.

The GitHub Copilot policy record was announced and configurable in September, with user-facing effect scheduled for October 22. The future effective date is explicit in the dataset.

The 20 verified changes

SEP-2026-001 / SHUTDOWN / EFFECTIVE 2026-09-01 / IMPACT HIGH

Six legacy Copilot model choices reached their September retirement

GitHub's dated retirement removed six model choices across most Copilot experiences. The annual individual-plan exception for Sonnet 4.6 means this was not a universal removal. Teams with pinned workflows need a supported replacement and an administrator policy check.

PREVIOUS STATE Six older model options remained selectable across most Copilot experiences.	NEW STATE Gemini 3.1 Pro, Claude Opus 4.5/4.6, Claude Sonnet 4.5/4.6 and Raptor Mini were deprecated across most Copilot experiences on September 1. Sonnet 4.6 remained available to individual annual subscribers.
WHO IS AFFECTED Copilot users and administrators with pinned models or model policies.	DECISION EFFECT Move pinned workflows to supported models and confirm enterprise model policies before migration.

SOURCE [GitHub changelog \(github.blog\)](#)

Announced or released 2026-09-01. Verified 2026-10-01.

SEP-2026-002 / MODEL OR PRODUCT LAUNCH / EFFECTIVE 2026-09-01 / IMPACT CRITICAL HIGH

Anthropic released Fable 5.1 with a restricted Mythos 5.1 variant

The shared-model relationship matters more than counting two unrelated launches. Fable 5.1 reached general use; Mythos 5.1 kept trusted-program access and different safeguards. Anthropic's cost claim depends on cache-heavy workload mix, not a blanket reduction to every token rate.

PREVIOUS STATE Fable 5 and Mythos 5 were the prior generation of Anthropic's highest-capability family.	NEW STATE Anthropic introduced Fable 5.1 for general use and Mythos 5.1 through trusted-access programs. The two share an underlying model but differ in safeguards; Anthropic lowered cache-read pricing while leaving standard token rates unchanged.
WHO IS AFFECTED Claude subscribers, API developers and vetted research organizations.	DECISION EFFECT Test real tasks at appropriate effort levels and check safeguard and access rules before assuming Mythos availability.

SOURCE [Anthropic announcement \(www.anthropic.com\)](#)

RELATED [Choosely AI Radar analysis](#)

Announced or released 2026-09-01. Verified 2026-10-01.

VERIFIED CHANGES / 03-04

SEP-2026-003 / MODEL OR PRODUCT LAUNCH / EFFECTIVE 2026-09-02 / IMPACT CRITICAL HIGH

Gemini 3.8 Flash arrived with a restricted Cyber counterpart

Google's third Flash release in six weeks strengthened the workhorse tier while holding its introductory token price. The Cyber variant shares foundational intelligence but has a different, restricted access path.

PREVIOUS STATE Gemini 3.7 Flash was Google's preceding low-cost Flash model.	NEW STATE Google introduced 3.8 Flash at an introductory US\$0.75 input and US\$3.75 output per million tokens, alongside Flash Cyber for trusted defenders in the Fairwind program.
WHO IS AFFECTED Gemini API builders, coding teams and vetted cybersecurity defenders.	DECISION EFFECT Re-evaluate Flash routing and costs; do not assume general access to the Cyber variant or permanent introductory pricing.

SOURCE [Google announcement \(blog.google\)](#)

Announced or released 2026-09-02. Verified 2026-10-01.

SEP-2026-004 / MODEL OR PRODUCT LAUNCH / EFFECTIVE 2026-09-03 / IMPACT CRITICAL HIGH

GPT-6 Astra made computer use and longer work a frontier buying decision

Astra's September 3 release materially expanded OpenAI's frontier tier for computer use, coding and professional work. The rollout was staged, and its highest-profile benchmark numbers depend on the harness used. The report treats provider claims as claims, not Choosely hands-on results.

PREVIOUS STATE GPT-5.6 Sol was the prior broadly deployed OpenAI frontier model.	NEW STATE OpenAI launched GPT-6 Astra with staged access across ChatGPT plans and the API; Enterprise administrators had to enable access. Standard API rates were US\$10 input and US\$50 output per million tokens.
WHO IS AFFECTED Professional teams, ChatGPT subscribers and API developers.	DECISION EFFECT Test complex end-to-end work with approval and security controls; distinguish provider benchmark settings from independent harness results.

SOURCE [OpenAI announcement \(openai.com\)](#)

RELATED [Choosely AI Radar analysis](#)

Announced or released 2026-09-03. Verified 2026-10-01.

VERIFIED CHANGES / 05-06

SEP-2026-005 / MODEL OR PRODUCT LAUNCH / EFFECTIVE 2026-09-09 / IMPACT HIGH

Suno v6 replaced the older generation across its music workflow

This is one model-family launch with a secondary retirement consequence, not a second counted shutdown. Creators should retest prompts and editing workflows because staying on the older generation is no longer the default option.

PREVIOUS STATE Suno offered older music models including v5.5.	NEW STATE Suno launched v6, v6-wild and v6-mini; the company says older models were retired. v6 and v6-wild require Pro or Premier, while v6-mini is available across plans.
WHO IS AFFECTED Music creators and teams with reproducible workflows built around earlier models.	DECISION EFFECT Archive reusable outputs and test editing, audio quality and plan access on the new family.

SOURCE [Suno announcement \(suno.com\)](#)

SOURCE [Suno v6 FAQ \(help.suno.com\)](#)

RELATED [Choosely AI Radar analysis](#)

Announced or released 2026-09-09. Verified 2026-10-01.

SEP-2026-006 / AVAILABILITY OR ACCESS CHANGE / EFFECTIVE 2026-09-14 / IMPACT HIGH

Siri AI shipped through a gradual, eligibility-gated beta

Installing iOS 27 did not guarantee the same experience for every user. Apple described a gradual English-language beta, supported-device requirements and regional exclusions, so availability is the primary September state change.

PREVIOUS STATE The announced next-generation personal Siri experience was not available in the public iOS release.	NEW STATE Apple released Siri AI with iOS 27 in a gradual English-language beta rollout, subject to supported hardware and regional requirements.
WHO IS AFFECTED iPhone users and organizations planning Apple Intelligence workflows.	DECISION EFFECT Check device, language, region and rollout eligibility before promising access; evaluate app actions on an actual supported device.

SOURCE [Apple announcement \(www.apple.com\)](#)

RELATED [Choosely AI Radar analysis](#)

Announced or released 2026-09-14. Verified 2026-10-01.

VERIFIED CHANGES / 07-08

SEP-2026-007 / CAPABILITY CHANGE / EFFECTIVE 2026-09-16 / IMPACT HIGH

Claude began folding Cowork into its main conversation

Anthropic removed a mode choice while preserving Cowork’s underlying capabilities. The rollout was gradual, so accounts could see the old and new interaction models during the transition.

PREVIOUS STATE Users chose between a separate Claude chat experience and Cowork mode for longer agentic tasks.	NEW STATE Anthropic began merging chat and Cowork into one Claude conversation, rolling out to Pro and Max over weeks; Claude Code remained separate.
WHO IS AFFECTED Claude Pro and Max users, with other plans to follow.	DECISION EFFECT Update workflow instructions and user training for the unified surface without describing Cowork’s capabilities as shut down.

SOURCE [Anthropic product announcement \(claude.com\)](#)

RELATED [Choosely AI Radar analysis](#)

Announced or released 2026-09-16. Verified 2026-10-01.

SEP-2026-008 / MODEL OR PRODUCT LAUNCH / EFFECTIVE 2026-09-16 / IMPACT HIGH

Claude added editable Docs and Slides in paid-plan beta

The same September 16 announcement also created distinct document and presentation products. This record counts those creation surfaces once; the Cowork integration is a separate change to an existing interaction model.

PREVIOUS STATE Claude did not offer the newly announced native Docs and Slides creation surfaces.	NEW STATE Anthropic introduced Claude Docs and Slides in beta on paid plans, with Enterprise administrator control; documents and presentations can be edited and exported.
WHO IS AFFECTED Paid Claude users and Enterprise administrators.	DECISION EFFECT Test export fidelity and admin settings on real deliverables before replacing document or slide workflows.

SOURCE [Anthropic product announcement \(claude.com\)](#)

RELATED [Choosely AI Radar analysis](#)

Announced or released 2026-09-16. Verified 2026-10-01.

VERIFIED CHANGES / 09-10

SEP-2026-009 / MODEL OR PRODUCT LAUNCH / EFFECTIVE 2026-09-17 / IMPACT HIGH

Gemini 3.8 Live split real-time voice work into two tiers

The models address different jobs: fluid voice at scale versus more reasoning-intensive multistep conversations. Google's quoted evaluation results remain vendor-reported evidence.

PREVIOUS STATE Gemini 3.8's newly announced live voice models were unavailable.	NEW STATE Google introduced 3.8 Live for scaled conversational use and 3.8 Live Extended Thinking for more complex multistep voice tasks.
WHO IS AFFECTED Developers and enterprises building voice agents.	DECISION EFFECT Compare latency, task success and running cost in the actual voice workflow before selecting a tier.

SOURCE [Google announcement \(blog.google\)](#)

Announced or released 2026-09-17. Verified 2026-10-01.

SEP-2026-010 / MODEL OR PRODUCT LAUNCH / EFFECTIVE 2026-09-21 / IMPACT HIGH

Grok 4.7 increased capability without changing its token rate

The vendor says Grok 4.7 checks work more carefully and handles longer jobs. Its rate card stayed level, but higher token consumption can still change cost per successful task.

PREVIOUS STATE Grok 4.6 was the preceding model at the published US\$2 input and US\$6 output rate per million tokens.	NEW STATE SpaceXAI released Grok 4.7 with a larger base model and longer training on hard tasks, at the same published token rate as 4.6.
WHO IS AFFECTED Coding teams and API developers using Grok.	DECISION EFFECT Compare completed-task cost and behavior against 4.6; the same token rate does not guarantee the same total spend.

SOURCE [SpaceXAI announcement \(x.ai\)](#)

RELATED [Choosely AI Radar analysis](#)

Announced or released 2026-09-21. Verified 2026-10-01.

VERIFIED CHANGES / 11-12

SEP-2026-011 / MODEL OR PRODUCT LAUNCH / EFFECTIVE 2026-09-22 / IMPACT CRITICAL HIGH

GPT-6 Sol and Luna reset the cost of OpenAI's working tiers

The paired release brought substantial API rate reductions alongside new model behavior. For Sol, published standard rates halved against GPT-5.6 promotional prices. Luna's output rate moved from US\$1.20 to US\$0.50; the vendor's broad '50% cheaper' language should not be applied to every token rate literally.

PREVIOUS STATE GPT-5.6 Sol and Luna were the previous models; their promotional rates were US\$4/\$20 and US\$0.20/\$1.20 input/output per million tokens.	NEW STATE OpenAI introduced GPT-6 Sol at US\$2/\$10 and Luna at US\$0.10/\$0.50 input/output per million tokens, alongside rollouts in Work, Codex and the API.
WHO IS AFFECTED API builders, ChatGPT Work and Codex users.	DECISION EFFECT Re-test task quality and effective cost by effort setting; do not count the associated price cuts as extra headline events.

SOURCE [OpenAI announcement \(openai.com\)](#)

RELATED [Choosely AI Radar analysis](#)

Announced or released 2026-09-22. Verified 2026-10-01.

SEP-2026-012 / MODEL OR PRODUCT LAUNCH / EFFECTIVE 2026-09-22 / IMPACT CRITICAL HIGH

Opus 5.5 raised Anthropic's frontier tier while lowering token rates

Anthropic quotes 40% lower typical-workload cost than Opus 5, while the standard input and output token rates each fell 20%. Those are different measurements. Subscription limit increases are a secondary consequence of this launch.

PREVIOUS STATE Opus 5 was the earlier Opus tier at US\$5 input and US\$25 output per million tokens.	NEW STATE Anthropic released Opus 5.5 at US\$4 input and US\$20 output, with US\$0.20 cache reads per million tokens; it also increased five-hour subscription limits.
WHO IS AFFECTED Claude subscribers and API teams routing harder work.	DECISION EFFECT Compare Opus 5.5 with Fable 5.1 and Sol on accepted-task cost, not nominal token price alone.

SOURCE [Anthropic announcement \(www.anthropic.com\)](#)

RELATED [Choosely AI Radar analysis](#)

Announced or released 2026-09-22. Verified 2026-10-01.

VERIFIED CHANGES / 13-14

SEP-2026-013 / MODEL OR PRODUCT LAUNCH / EFFECTIVE 2026-09-23 / IMPACT HIGH

Gemini 3.8 added a directed text-to-speech model family

The launch covered both full Flash TTS and Flash-Lite TTS and introduced more direct control of voices and dialogue. It is counted once as a related model-family release.

PREVIOUS STATE Gemini's newly announced 3.8 text-to-speech family was not available.	NEW STATE Google introduced 3.8 Flash TTS and Flash-Lite TTS across AI Studio, the API and selected enterprise/content surfaces, including voice direction and character creation.
WHO IS AFFECTED Audio creators and developers building speech workflows.	DECISION EFFECT Check generated speech quality, permitted use and platform-specific access before changing production audio workflows.

SOURCE [Google announcement \(blog.google\)](#)

Announced or released 2026-09-23. Verified 2026-10-01.

SEP-2026-014 / POLICY CHANGE / EFFECTIVE 2026-10-22 / IMPACT HIGH

Copilot admins gained a future default policy with an October action date

The September event is the availability of a policy control and a published 28-day transition, not a claim that user access already changed. The effective change is scheduled for October 22.

PREVIOUS STATE Unconfigured Copilot feature policies did not follow the newly announced global default.	NEW STATE Administrators could configure a global default from September 24; it is scheduled to govern eligible unconfigured features beginning October 22. Explicit decisions and preview opt-ins remain preserved.
WHO IS AFFECTED Copilot Business and Enterprise administrators.	DECISION EFFECT Set the default deliberately before October 22 and review unconfigured feature, code-review and MCP policies.

SOURCE [GitHub changelog \(github.blog\)](#)

Announced or released 2026-09-24. Verified 2026-10-01.

VERIFIED CHANGES / 15-16

SEP-2026-015 / MODEL OR PRODUCT LAUNCH / EFFECTIVE 2026-09-24 / IMPACT HIGH

Gemini 3.8 Live Avatar put a video face on enterprise voice agents

This is a distinct enterprise visual surface built on the September 17 live-voice models. It is counted separately because it changes the interface and buying decision, while the underlying voice release is not counted again.

PREVIOUS STATE Gemini 3.8 Live voice models lacked the newly announced synchronized live-avatar output.	NEW STATE Google made Live Avatar available in Gemini Enterprise, combining real-time dialogue with synchronized generated video and background tool calls.
WHO IS AFFECTED Enterprises building short interactive service and guided-support agents.	DECISION EFFECT Pilot duration, latency, disclosure and human handoff before promising an always-on digital host.

SOURCE [Google announcement \(blog.google\)](#)

RELATED [Choosely AI Radar analysis](#)

Announced or released 2026-09-24. Verified 2026-10-01.

SEP-2026-016 / MODEL OR PRODUCT LAUNCH / EFFECTIVE 2026-09-28 / IMPACT HIGH

Sonnet 5.5 sharpened the price and effort choice inside Claude

The September 28 launch completed a new Sonnet/Opus comparison in the 5.5 family. Half-price standard tokens do not guarantee half the cost per completed task at high effort.

PREVIOUS STATE Sonnet 5 was the prior Sonnet tier.	NEW STATE Anthropic released Sonnet 5.5 as a faster, lower-cost complement to Opus 5.5, with standard API input and output rates half the Opus 5.5 rates.
WHO IS AFFECTED Claude users, API builders and teams routing volume workloads.	DECISION EFFECT Use Sonnet for well-scoped work, then compare Opus as effort and total task cost rise.

SOURCE [Anthropic announcement \(www.anthropic.com\)](#)

RELATED [Choosely AI Radar analysis](#)

Announced or released 2026-09-28. Verified 2026-10-01.

VERIFIED CHANGES / 17-18

SEP-2026-017 / MODEL OR PRODUCT LAUNCH / EFFECTIVE 2026-09-29 / IMPACT CRITICAL HIGH

GPT-6.1 Sol arrived one week after GPT-6 Sol

OpenAI's September 29 changelog lists the new model in Responses and Chat Completions, including beta multi-agent support. Its rapid follow-on timing raises a practical version-pinning and retesting question.

PREVIOUS STATE GPT-6 Sol was the current work and coding tier.	NEW STATE OpenAI released GPT-6.1 Sol for coding and professional work at standard US\$2 input and US\$10 output per million tokens, with staged Work and Codex access.
WHO IS AFFECTED API developers and users of Work and Codex.	DECISION EFFECT Retest coding and computer-use workflows; do not confuse released 6.1 Sol with unreleased 6.1 Astra.

SOURCE [OpenAI API changelog \(developers.openai.com\)](#)

Announced or released 2026-09-29. Verified 2026-10-01.

SEP-2026-018 / MODEL OR PRODUCT LAUNCH / EFFECTIVE 2026-09-29 / IMPACT CRITICAL HIGH

ChatGPT dots introduced persistent delegated work

Dots can continue work between conversations on an assigned cloud computer and return results for review. Access was gradual and geographically limited, while introductory usage treatment was temporary.

PREVIOUS STATE ChatGPT did not provide the newly announced persistent dot agent product.	NEW STATE OpenAI began a gradual adult-user rollout of Astra-powered dots to eligible Pro and Business Premium users in supported markets; Enterprise beta is off by default.
WHO IS AFFECTED Eligible Pro, Business Premium and Enterprise users administering ongoing work.	DECISION EFFECT Start with bounded tasks and explicit app permissions; check geography, plan and post-introductory usage terms.

SOURCE [OpenAI Help Center release notes \(help.openai.com\)](#)

Announced or released 2026-09-29. Verified 2026-10-01.

VERIFIED CHANGES / 19-20

SEP-2026-019 / PLAN CHANGE / EFFECTIVE 2026-09-29 / IMPACT HIGH

Pro 500 made Astra Ultrafast a separate subscription decision

The new tier changed plan inclusion and access rather than simply repricing an existing plan. OpenAI lists US\$500 per month and says Ultrafast is exclusive to Pro 500 among Pro plans.

PREVIOUS STATE Pro 100 and Pro 200 did not include access to Astra Ultrafast in Work and Codex.	NEW STATE OpenAI introduced a US\$500/month Pro 500 tier with the highest included Pro usage and Astra Ultrafast access; buying credits on lower Pro tiers does not unlock Ultrafast.
WHO IS AFFECTED Heavy individual Work and Codex users comparing subscriptions.	DECISION EFFECT Compare actual latency and allowance use before upgrading; do not treat top-up credits as a substitute for plan eligibility.

SOURCE [OpenAI Help Center release notes \(help.openai.com\)](#)

Announced or released 2026-09-29. Verified 2026-10-01.

SEP-2026-020 / CAPABILITY CHANGE / EFFECTIVE 2026-09-29 / IMPACT HIGH

The Agents API gained hosted-browser computer use

This is a developer API capability separate from the consumer-facing dots product. It changes the deployment and control path for browser-acting agents.

PREVIOUS STATE The Agents API lacked the newly announced OpenAI-hosted browser computer-use capability.	NEW STATE OpenAI added hosted-browser computer use with website access approvals and application-managed sign-in.
WHO IS AFFECTED Developers building agents that act on websites.	DECISION EFFECT Design approval boundaries, login handling and audit trails before enabling web actions.

SOURCE [OpenAI API changelog \(developers.openai.com\)](#)

Announced or released 2026-09-29. Verified 2026-10-01.

Pricing and plan analysis

Token rate is not task cost

Sol's rate reductions and Opus's cheaper cache reads can be material, but effort, output length and repeated context reads determine actual spend.

Plan inclusion is its own event

Pro 500's Ultrafast entitlement is not unlocked by buying credits on Pro 100 or Pro 200.

Introductory rates have an end

Google's Gemini 3.8 Flash introductory rate should be modeled separately from a permanent price assumption.

Category analysis

Frontier capability became a routing problem

Astra, Fable 5.1, Opus 5.5, Gemini 3.8 Flash, Grok 4.7, Sonnet 5.5 and two Sol releases landed within one month. The vendor materials describe meaningful gains, but their benchmark harnesses, effort settings and safeguards differ. A raw score is not a universal buying order.

For teams, the durable comparison is the cost and reliability of a completed job under the actual plan, tool permissions and approval process. The one-week gap between GPT-6 Sol and GPT-6.1 Sol makes version pinning and repeatable evaluations especially practical.

The agent surface moved into ordinary work

The Claude integration, native Docs and Slides, ChatGPT dots and hosted browser computer use make multi-step work more accessible. They also cross different boundaries: an interactive conversation, a persistent cloud computer and a developer-controlled API are not interchangeable deployment models.

Give these systems narrow permissions, review the actions they take and retain a fallback. The report does not assume Choose.ly tested every product hands-on.

Voice and video became separate purchasing decisions

Gemini 3.8 Live, its Extended Thinking tier, the TTS pair and Live Avatar made voice and visual presence more specialized. A conversational model, directed speech generator and animated enterprise agent solve different problems even when they share a family name.

Evaluate latency, handoff, disclosure, platform access and full-session cost. The short enterprise avatar interaction documented by Google should not be sold as an unlimited digital presenter.

What to watch in October

The October watchlist is not included in September event totals.

- GitHub's October 22 Copilot default-policy effect and any actual administrator/user access consequences.
- Rollout and post-introductory terms for ChatGPT dots and Astra Ultrafast.
- Evidence of Gemini Flash introductory pricing after December 31 and any confirmed replacement terms.
- Further model retirements only when effective, not merely announced for a future month.

Methodology

Scope and dates

The reporting window is September 1 through September 30, 2026, with first-party evidence checked through October 1. A release or effective state change qualifies when it materially changes a user's access, workflow, cost, governance or migration decision. One September announcement with an October effective date is clearly labeled as a scheduled transition.

A common announcement can contain separate records only when it changes distinct products or decisions. Related variants within one model family are grouped. One primary class is assigned to each record; secondary price or retirement consequences do not add counts.

Evidence and record fields

Official vendor announcements, changelogs, pricing, support and policy documents carry factual claims. Choosely AI Radar helps interpret the change but does not independently verify it. Every public record has a stable SEP-2026-NNN ID, a primary class, date, source, affected users and a practical decision effect.

The typed dataset generates the CSV. The HTML is canonical, and the PDF is a readable supporting edition of the same 20 records. A prior state is stated only where the source or prior record supports it.

Limitations

- This is a selected verified dataset, not a census of every vendor, beta, region, enterprise contract or catalog item. A zero-count class is not evidence that no such changes occurred anywhere.
- Vendor benchmark and speed claims are attributed to vendors. This report does not claim a controlled Choosely test of the listed models.
- Rollouts, eligibility and rate cards can change after the October 1 verification cutoff. Check current vendor terms before buying or migrating.
- The dataset omits private monitoring mechanics, internal candidate notes and personal account data.

Downloadable data and reuse terms

Canonical HTML: <https://choosely.ai/reports/ai-tool-changes/2026-09>

PDF: </reports/ai-tool-changes/2026-09/report.pdf>

CSV: </reports/ai-tool-changes/2026-09/data.csv>

The HTML is canonical; this PDF and the generated CSV contain the same 20 event-level records.

© 2026 Choosely. All rights reserved. The dataset may be quoted and referenced with attribution. Redistribution, republication or commercial reuse of the complete dataset requires written permission from Choosely.

Cite this report

Choosely Editorial. 20 Verified AI-Tool Changes That Mattered in September 2026. Choosely, October 1, 2026.
<https://choosely.ai/reports/ai-tool-changes/2026-09>

Corrections and version history

Material corrections retain stable event IDs and are dated across HTML, PDF and CSV. Factual correction requests: info@choosely.ai.

Version 1.0 published October 1, 2026.

Final September verdict

September brought more capable models, but the operator's real decision moved toward routing, access and control. Choose the model and surface that completes your work at an acceptable cost, then verify what it can reach and what a human must approve.

Primary source register

First-party evidence for the 20 locked September records. Full links are clickable and retained in the CSV.

SEP-2026-001 <https://github.blog/changelog/2026-08-31-selected-github-copilot-models-deprecated/>
SEP-2026-002 <https://www.anthropic.com/claude-fable-and-mythos-5-1>
SEP-2026-003 <https://blog.google/innovation-and-ai/models-and-research/gemini-models/3-8-flash-and-3-8-flash-cyber/>
SEP-2026-004 <https://openai.com/index/gpt-6-astra/>
SEP-2026-005 <https://suno.com/blog/introducing-v6>
SEP-2026-005 <https://help.suno.com/en/articles/13924481>
SEP-2026-006 <https://www.apple.com/newsroom/2026/09/siri-ai-a-profoundly-more-capable-and-personal-assistant-is-here/>
SEP-2026-007 <https://claude.com/blog/cowork-is-now-claude>
SEP-2026-008 <https://claude.com/blog/cowork-is-now-claude>
SEP-2026-009 <https://blog.google/innovation-and-ai/models-and-research/gemini-models/gemini-3-8-live-gemini-3-8-live-extended-thinking/>
SEP-2026-010 <https://x.ai/news/grok-4-7>
SEP-2026-011 <https://openai.com/index/introducing-gpt-6-sol-and-luna/>
SEP-2026-012 <https://www.anthropic.com/claude-opus-5-5>
SEP-2026-013 <https://blog.google/innovation-and-ai/models-and-research/gemini-models/gemini-3-8-text-to-speech/>
SEP-2026-014 <https://github.blog/changelog/2026-09-24-default-enablement-of-copilot-features-for-copilot-business-and-enterprise/>
SEP-2026-015 <https://blog.google/innovation-and-ai/models-and-research/gemini-models/gemini-3-8-live-with-live-avatar/>
SEP-2026-016 <https://www.anthropic.com/claude-sonnet-5-5>
SEP-2026-017 <https://developers.openai.com/api/docs/changelog>
SEP-2026-018 <https://help.openai.com/en/articles/6825453-chatgpt-release-notes>
SEP-2026-019 <https://help.openai.com/en/articles/6825453-chatgpt-release-notes>
SEP-2026-020 <https://developers.openai.com/api/docs/changelog>