FOR YOU / DAILY

AI Daily

2026-11-10 • Cutoff 2026-11-10T08:30:00+08:00

Highlights

Claude 5.2 launches with two access tiers

Fable serves general users; Mythos requires specialist verification. Standard API input/output costs $8/$40 per million tokens.

Anthropic: Claude Fable and Mythos 5.2

Qwen Office 2.0 adds triggers and mobile access

Team spaces connect DingTalk work data and trigger digital employees from events; account, organization and administrator controls constrain remote access.

Qwen Office desktop 2.0.0 release notes

The Beijing News: Qwen Office 2.0

Google announces TPU 8tx and 8i Max

8tx changes hosts and networking; 8i Max packages two 8i chips. Application throughput and general availability remain separate questions.

Google: TPU 8i Max and 8tx

01 / AI companies

2026-11-09 • New

Anthropic launches Claude Fable and Mythos 5.2 with different access requirements

Anthropic released Fable and Mythos 5.2 using the same underlying model. Fable serves general users; Mythos requires its cybersecurity or life-sciences verification program. Standard API input/output pricing is $8/$40 per million tokens. Fable applies cyber and bio safeguards; teams must choose access appropriate to their tasks and eligibility. The optional faster mode costs more.

Anthropic: Claude Fable and Mythos 5.2

2026-11-09 • New

ByteDance releases Seedream 5.5 with iterative editing and transparent images

ByteDance Seed released Seedream 5.5, demonstrating consistency across 100 successive edits and native RGBA transparency. Its official page links to Dola and BytePlus Playground. Transparency and localized editing help create reusable design assets. The 100-round result is a company demonstration, not a guarantee for every image. The higher-specification Pro version is still forthcoming.

ByteDance Seed: Seedream 5.5

Seedream 5.5 official access

2026-11-09 • New

Qwen Office 2.0 adds team spaces, event triggers and mobile remote control

Qwen Office desktop 2.0.0 adds team spaces connected to DingTalk documents, AI tables and Yida. Messages, approvals and tasks can trigger digital employees. The mobile app is released with remote control of enabled computers. Teams can turn one-off chats into event-triggered workflows. Existing document permissions remain in force; remote control requires the same account and organization, and administrators can disable it.

Qwen Office desktop 2.0.0 release notes

The Beijing News: Qwen Office 2.0

2026-11-09 • New

Baidu presents ERNIE Evolve for automated model-engineering experiments

Baidu presented ERNIE Evolve, combining isolated experiments with a separate evaluation stage, and released selected code and logs. It reports performance reaching 6.236 times baseline after four optimization rounds in a low-latency Kunlunxin P800 setting. The records let developers inspect changes and evaluation procedures. The figure is Baidu’ s configuration-specific result, not external replication. This is an automated research method, not a new launch of the model under development.

Baidu: ERNIE Evolve automated research

ERNIE Evolve experiment code and records

2026-11-09 • New

OpenAI deploys Control Room to centrally stop research-model outputs

OpenAI described Control Room, deployed across its research-inference infrastructure for agent training, evaluations and data processing. It can stop model outputs within 35 seconds of an intervention decision. Central control provides a stop mechanism across research clusters and fails closed if the control service fails. The 35 seconds measures post-decision output cessation, not risk-detection time or termination of every underlying workload.

OpenAI: Control Room

2026-11-07 • Seven-day watch

Google researchers use Lean to verify AI-generated quantum programs

Google DeepMind and Quantum AI researchers published a method in which model agents propose quantum programs against Lean specifications and verify correctness and resource costs, collaborating through shared knowledge repositories. The approach offers checkable proofs for quantum algorithms that are difficult to simulate directly. Reported improvements concern algorithms and cost models, not measured hardware speedups. The original November 6 UTC submission falls on November 7 in Beijing.

Discovering formally verified quantum programs with AI

Quantum-program paper affiliations and method

02 / AI product tracker

New this issue

Claude Fable · Anthropic

5.2 • 2026-11-09 • New this issue • stable

This issue: general-user access; standard API input/output pricing is $8/$40 per million tokens.

Anthropic: Claude Fable and Mythos 5.2

Claude Mythos · Anthropic

5.2 • 2026-11-09 • New this issue • stable

This issue: verified-access variant of the same model; cybersecurity or life-sciences verification is required.

Anthropic: Claude Fable and Mythos 5.2

Seedream · ByteDance

5.5 • 2026-11-09 • New this issue • stable

This issue: iterative editing and RGBA transparency; Dola and BytePlus Playground links are offered. Pro is forthcoming.

ByteDance Seed: Seedream 5.5

Seedream 5.5 official access

Qwen Office · Alibaba

2.0.0 • 2026-11-09 • New this issue • stable

This issue: team spaces, event triggers and mobile remote control, subject to account, organization and administrator permissions.

Qwen Office desktop 2.0.0 release notes

Recently verified versions

Bailian CLI · Alibaba Cloud

2.8.4 • 2026-11-06 • Recent version • stable

Recent verified baseline: fine-tuning parameter and metrics updates; not a new release in this issue.

Bailian CLI release notes

03 / AI chips

2026-11-09 • New

Google announces TPU 8tx and 8i Max with new hosts, networking and packaging

Google announced two TPU system configurations. 8tx pairs 8t silicon with fifth-generation AMD EPYC hosts and Boardfly networking. 8i Max packages two 8i chips together, doubling compute, SRAM and HBM bandwidth versus one 8i, according to Google. 8tx targets confidential inference and private fine-tuning; 8i Max targets very large context windows. No general-availability date was established in the announcement reviewed here, and doubled chip resources do not guarantee doubled application throughput.

Google: TPU 8i Max and 8tx

04 / Embodied AI

2026-11-06 • Seven-day watch

Meituan shares MR0 robot research and training code across 12 embodiments

Meituan’s robotics team shared MR0 vision-language-action research with training and evaluation code. The paper uses 147 million vision-language samples and 15,060 hours of robot trajectories spanning 12 embodiments. Physical demonstrations involving blocks, clothes and cables distinguish it from digital-only agents. Code and evaluation interfaces support scrutiny, but these remain research results and demonstrations, not deliveries of general-purpose robots or unattended production capability.

Meituan Robotics 0 paper

Meituan MR0 training and evaluation code

05 / Infrastructure

2026-11-08 • Recent, within 72 hours

NVIDIA, Rebellions and SK Telecom plan hybrid inference infrastructure; R200 targets 2028

NVIDIA, Rebellions and SK Telecom announced a collaboration using Rebellions R200 NPUs, NVLink Fusion and MGX. The plan assigns prefill to NVIDIA and decode to Rebellions, with SK Telecom pursuing data-center deployment. The design assigns inference stages to different accelerators, offering an integration path for Korean infrastructure. R200 systems are expected to be deployment-ready in 2028; integration, delivery and real-workload efficiency remain to be demonstrated.

NVIDIA, Rebellions and SK Telecom collaboration

06 / Policy

2026-11-09 • New

Six countries and ten AI companies sign Paris declaration on evaluations and risk reporting

Six countries including Canada and France, ten AI companies and four civil-society organizations signed the Paris AI safety declaration, committing to shared risk methods, third-party evaluations, incident information sharing and a repository of non-sensitive findings. Company commitments also address fail-safe and shutdown mechanisms for training and internal models. Developers now need to turn commitments into inspectable evaluation and reporting processes. The declaration is not a newly binding law or proof that the repository already operates.

Canadian PM: Paris AI safety and security declaration

07 / Business

2026-11-09 • New

Adobe and Disney expand multiyear alliance for custom Firefly creative models

Adobe and Disney expanded their multiyear strategic alliance, bringing Firefly Foundry custom models into creative and experience development and involving StudioLAB and ESPN Edge. The companies had already worked together in Imagineering. The alliance combines proprietary content with engineering support for workflows such as concept design. Some projects are under way and further activations are forthcoming. No financial value was disclosed, and the custom models are not thereby available to the public.

Adobe and Disney expand multiyear alliance

2026-11-08 • Recent, within 72 hours

NVIDIA board opens independent inquiry into reports Huang was targeted by foreign intelligence

NVIDIA’s board announced an inquiry by independent directors and outside counsel into reports that Jensen Huang was targeted by foreign intelligence. It says it has no reason to believe national or company secrets were compromised, and Huang will cooperate. The development is the opening of a corporate-governance inquiry. Its scope and formal findings are the next disclosures to watch. The statement neither establishes a leak nor turns reported allegations into proven facts.

NVIDIA board announces independent inquiry

2026-11-06 • Seven-day watch

Phison reports NT$86.076 billion Q3 revenue and an AI-ecosystem share above half

Phison reported third-quarter consolidated revenue of NT$86.076 billion, up 374.6% year on year, and a 59.7% gross margin. Its broadly defined AI ecosystem accounted for more than half of revenue, including enterprise SSDs, AI PCs, networking and servers. This is realized storage-supply-chain revenue, rather than a funding valuation or order target; the broad category is not pure AI-model revenue. Inventory turnover days rose from 224 a year earlier to 274, making subsequent shipment absorption relevant.

Phison third-quarter 2026 results

08 / AI markets

Most US watchlist names rise while NVIDIA falls; mainland and Hong Kong names diverge

On November 9, the US watchlist had 14 gainers and one decliner: AMD rose 13.26%, Western Digital 9.86%, while NVIDIA fell 4.27%. Mainland names split 3 up/5 down, with Cambricon +6.02% and CXMT -1.04%; Hong Kong split 2 up/3 down. SK hynix rose 0.35% and Samsung Electronics 0.64%.

These 30 companies span five industry groups and do not form an industry index. Market closes are asynchronous. USD, CNY, HKD and KRW are retained; prices and changes use each market's latest complete session, with no missing values today. Source: limited excerpts from Stock Analysis history. Moves are observations, not news-based causal claims.

Memory

Company / ticker	Market / FX	Close	Change	Session / source
Micron / MU	US / USD	733.15	+6.25%	2026-11-09 Micron November 9 close
Sandisk / SNDK	US / USD	886.51	+6.16%	2026-11-09 Sandisk November 9 close
SK hynix / 000660	KR / KRW	1,136,000	+0.35%	2026-11-09 SK hynix November 9 close
Samsung Electronics / 005930	KR / KRW	189,900	+0.64%	2026-11-09 Samsung Electronics November 9 close
CXMT / 688825	CN / CNY	39.99	-1.04%	2026-11-09 CXMT November 9 close

Storage and components

Company / ticker	Market / FX	Close	Change	Session / source
Western Digital / WDC	US / USD	289.78	+9.86%	2026-11-09 Western Digital November 9 close
Seagate / STX	US / USD	506.14	+9.79%	2026-11-09 Seagate November 9 close
Montage Technology / 688008	CN / CNY	149.67	-0.29%	2026-11-09 Montage Technology November 9 close
GigaDevice / 603986	CN / CNY	238.41	-0.45%	2026-11-09 GigaDevice November 9 close
Longsys / 301308	CN / CNY	215.56	+0.37%	2026-11-09 Longsys November 9 close
BIWIN / 688525	CN / CNY	125.41	-0.47%	2026-11-09 BIWIN November 9 close

Compute and manufacturing

Company / ticker	Market / FX	Close	Change	Session / source
Broadcom / AVGO	US / USD	319.48	+8.78%	2026-11-09 Broadcom November 9 close
TSMC ADR / TSM	US / USD	388.62	+3.51%	2026-11-09 TSMC ADR November 9 close
NVIDIA / NVDA	US / USD	178.67	-4.27%	2026-11-09 NVIDIA November 9 close
AMD / AMD	US / USD	508.58	+13.26%	2026-11-09 AMD November 9 close
Intel / INTC	US / USD	59.37	+5.49%	2026-11-09 Intel November 9 close
Cambricon / 688256	CN / CNY	851.21	+6.02%	2026-11-09 Cambricon November 9 close
Hygon / 688041	CN / CNY	191.88	+1.26%	2026-11-09 Hygon November 9 close
SMIC / 00981	HK / HKD	46.04	+1.01%	2026-11-09 SMIC November 9 close

Networking, power and cooling

Company / ticker	Market / FX	Close	Change	Session / source
Arista Networks / ANET	US / USD	186.17	+9.00%	2026-11-09 Arista Networks November 9 close
Vertiv / VRT	US / USD	200.47	+9.63%	2026-11-09 Vertiv November 9 close

Platforms and applications

Company / ticker	Market / FX	Close	Change	Session / source
Microsoft / MSFT	US / USD	545.18	+2.75%	2026-11-09 Microsoft November 9 close
Alphabet / GOOGL	US / USD	288.15	+1.18%	2026-11-09 Alphabet November 9 close
Meta / META	US / USD	641.37	+5.12%	2026-11-09 Meta November 9 close
Amazon / AMZN	US / USD	266.93	+2.94%	2026-11-09 Amazon November 9 close
iFlytek / 002230	CN / CNY	37.96	-0.76%	2026-11-09 iFlytek November 9 close
Baidu / 09888	HK / HKD	78.50	+0.38%	2026-11-09 Baidu November 9 close
Alibaba / 09988	HK / HKD	104.60	-0.29%	2026-11-09 Alibaba November 9 close
Tencent / 00700	HK / HKD	418.40	-0.33%	2026-11-09 Tencent November 9 close
SenseTime / 00020	HK / HKD	1.15	-1.29%	2026-11-09 SenseTime November 9 close

Next, watch Mythos verification access, Qwen team-permission deployment, TPU availability plans and implementation of Paris evaluation and reporting commitments.