
FOR YOU / HIGHLIGHTS • EDITOR'S LETTER

Models advance. The systems around them must follow.

Kimi brings a million-token-context model to products and APIs, Agent Skills moves toward neutral governance, and China publishes AI employment guidelines. Today we connect model capability with skills standards, packaging capacity, data and operations: what makes the new capability usable?

Model availability, reliable systems and workforce transitions belong in the same deployment plan.

Kimi enters more workflows

Million-token context is available; weights are promised by November 4.

[Kimi: K3.1 Pro release and availability](#)

[Kimi official blog dates](#)

[Kimi Code release notes](#)

Put workforce transitions in the AI plan

New guidelines address job-risk assessment and transition support.

[State Council: Employment-friendly AI development](#)

Capacity needs remain; funding paths change

Solidigm ends its IPO review; alternative funding remains undecided.

[SK hynix: Solidigm discontinues IPO review](#)

[SK hynix dated news index](#)

01 / AI companies

2026-10-24 • New

Kimi K3.1 Pro goes live; hosted access precedes weights

Kimi released K3.1 Pro and Highspeed with a one-million-token context window across the web, mobile apps, Kimi Work, Kimi Code and API. Full model weights are promised by November 4; enterprise administrators must enable access, which is off by default at launch.

Teams can now test long-document and cross-file tasks, while self-hosting still awaits weights. Performance and speed figures are vendor-reported; assess task completion, context use and total cost on your own workloads. API availability and weight availability are separate milestones.

[Kimi: K3.1 Pro release and availability](#)

[Kimi official blog dates](#)

[Kimi Code release notes](#)

2026-10-22 • Recent

Agent Skills moves to AAIF for neutral cross-tool governance

Anthropic announced the donation of the Agent Skills name, logo, GitHub organization and domain to AAIF under the Linux Foundation. It says more than 70 agent platforms have integrated the standard; using or contributing to it does not require foundation membership.

Neutral governance can support continued reuse of internal skills across tools. Versioned releases, a formal proposal process and security guidance remain a proposed roadmap; transferring the standard does not certify the safety of individual skills.

[Anthropic: Agent Skills donation to AAIF](#)

2026-10-23 • New

Codex 0.164.0 brings reusable prompts and enterprise access into terminal workflows

The stable release adds ten numbered /presets prompt slots, inline model and permission controls, and enterprise MCP authorization improvements. Release notes also describe credential masking for supported remote Bash/Zsh sessions and state fixes after failed conversation compaction.

The value is less repeated setup and easier recovery during long tasks. Credential masking covers supported paths, not a guarantee that every log is free of sensitive data. Teams should verify their own tool permissions and recovery workflows.

[Codex 0.164.0 release notes](#)

02 / AI product tracker

Product	Version / date	Changes and sources
Kimi	K3.1 Pro / Highspeed 2026-10-24 New this issue • stable	Hosted release is available with a million-token context; enterprise access is off by default and can be enabled by administrators. Weights are still promised by November 4. Kimi: K3.1 Pro release and availability Kimi official blog dates
Codex	0.164.0 2026-10-23 New this issue • stable	Stable: Prompt presets, inline controls and enterprise MCP authorization improvements. See AI Companies. Codex 0.164.0 release notes
Claude Code	2.1.310 2026-10-24 New this issue • stable	Stable: Extends Monitor availability and fixes symlink permission prompts, a Windows hook-approval path and unintended plan approval from VS Code feedback. Claude Code 2.1.310 release notes

03 / AI chips

2026-10-22 • Recent

JCET reports 54% computing-electronics growth; capacity still needs delivery

JCET reports third-quarter revenue of about CNY11.425 billion, up 13.53% year on year. Computing electronics represented 30.9% of nine-month revenue and grew 54%. Its Lingang advanced-packaging project is progressing, while related expansion financing still requires approvals.

These operating figures help track compute demand reaching packaging, but computing-electronics revenue is not all AI revenue. The quarterly report is unaudited; the project and financing plans began in June and September. This item covers the new operating disclosure, not planned capacity as already operational.

[JCET third-quarter 2026 report](#)

04 / Embodied AI

2026-10-21 • Recent

Cosmos expands five data blueprints for hard-to-capture robot scenarios

Cosmos Data Factory offers five blueprints for driving datasets, image attributes, event videos, defect images and video augmentation. NVIDIA reports that ANYbotics generated more than 1,500 synthetic clips to help industrial-inspection robots recognize liquid spills and other scene changes.

The aim is better coverage of rare events, not reliable robots simply from more videos. Generated outputs still need validation, real-data integration and task-level evaluation. Projected testing and development savings are not independently reproduced outcomes.

[NVIDIA: Cosmos data and model tools](#)

05 / Infrastructure

2026-10-22 • Recent

Minerva connects GPUs, networks and jobs to locate bottlenecks

NVIDIA launched Minerva, an AIOps platform connecting GPU, network, storage, scheduler and workload telemetry in a dynamic knowledge graph to correlate alerts, detect anomalies and identify likely root causes. IREN is using it; partners including IBM and Lambda plan deployments.

Beyond buying compute, operators need to measure useful work and recovery time. Minerva can recommend remediation or connect to automation, while teams can keep execution with people. Recovery and throughput benefits still need validation on the actual cluster.

[NVIDIA: Minerva operations platform](#)

2026-10-23 • New

OpenAI storage lessons: controlled sharing can beat more isolation

OpenAI describes Habitat moving from shared to dedicated containers, then back to sharing with hot-key throttling, request priorities and improved serverless partition scaling. Dedicated containers reduced failure spread but pushed capacity planning onto individual product teams.

For large AI services, data-layer operational complexity can slow product iteration. The article also discusses caching and cross-region capacity-versus-latency tradeoffs. A next-generation system based on RocksDB Cloud and Raft is still being developed, not a completed replacement.

[OpenAI: Storage at one-billion-user scale, part two](#)

06 / Policy

2026-10-23 • New

China brings employment effects into AI project assessment

China's State Council General Office published employment-friendly AI guidelines. They call for employment-risk assessment alongside review of market-entry projects that may cause substantial job displacement, and encourage employers to explore transition buffers, advance notice and redeployment support.

An AI deployment plan should therefore include training, job transitions and worker protections alongside procurement and efficiency targets. The document was dated October 16 and published October 23. Encouraging transition buffers does not impose one immediate, uniform duration on every employer.

[State Council: Employment-friendly AI development](#)

2026-10-22 • Recent

US proposes broader outbound-investment rules for AI and compute

The US Treasury proposed COINS implementing rules covering technologies including semiconductors, AI, quantum and high-performance computing, with biotechnology also proposed. The framework distinguishes notifiable and prohibited transactions and addresses covered persons, exceptions and reasonable pre-transaction inquiry.

Cross-border investors need to examine technical parameters, transaction structure and covered persons, rather than a blanket notion of an AI investment ban. This remains a proposal for comment; final regulations and their effective date are still to follow.

[US Treasury: Proposed COINS implementing rules](#)

07 / Business

2026-10-23 • New

Solidigm ends its IPO review while considering other funding routes

SK hynix announced that Solidigm will stop reviewing potential pre-IPO and IPO options, citing capital-market conditions and share-price trends. It still sees a need for additional US production facilities; investment needs are unchanged and alternative funding arrangements remain under consideration.

Memory demand and financing routes can move at different speeds. Ending a listing review does not cancel factory investment or mean a filed IPO was withdrawn. The announcement confirms neither alternative funding amounts nor commissioning dates or new available capacity.

[SK hynix: Solidigm discontinues IPO review](#)

[SK hynix dated news index](#)

08 / AI markets

Asian watchlist gains; US storage names diverge

At the October 23 closes, 21 watchlist companies rose and nine fell. Mainland CXMT gained 7.93% and GigaDevice 5.51%; SK hynix gained 4.42% in Korea. US-listed Micron rose 2.46%, while Sandisk, Western Digital and Seagate fell 1.82%, 2.97% and 1.81%.

Market closes are not simultaneous; the watchlist is not an industry index. Original USD/CNY/HKD/KRW currencies are retained. All 30 prices and changes are sourced, with no individual missing values. Later news is not used to explain earlier prices. Limited attributed excerpts from Stock Analysis.

Memory

Company / ticker	Market / FX	Close	Change	Session / source
Micron / MU	US / USD	877.55	+2.46%	2026-10-23 Stock Analysis • MU
Sandisk / SNDK	US / USD	1,238.03	-1.82%	2026-10-23 Stock Analysis • SNDK
SK hynix / 000660	KR / KRW	1,277,000	+4.42%	2026-10-23 Stock Analysis • 000660
Samsung Electronics / 005930	KR / KRW	215,500	+0.47%	2026-10-23 Stock Analysis • 005930
CXMT / 688825	CN / CNY	42.85	+7.93%	2026-10-23 Stock Analysis • 688825

Storage and components

Company / ticker	Market / FX	Close	Change	Session / source
Western Digital / WDC	US / USD	343.67	-2.97%	2026-10-23 Stock Analysis • WDC
Seagate / STX	US / USD	630.49	-1.81%	2026-10-23 Stock Analysis • STX
Montage Technology / 688008	CN / CNY	180.00	+2.78%	2026-10-23 Stock Analysis • 688008
GigaDevice / 603986	CN / CNY	292.00	+5.51%	2026-10-23 Stock Analysis • 603986
Longsys / 301308	CN / CNY	244.68	+1.11%	2026-10-23 Stock Analysis • 301308
BIWIN / 688525	CN / CNY	164.97	+3.43%	2026-10-23 Stock Analysis • 688525

Compute and manufacturing

Company / ticker	Market / FX	Close	Change	Session / source
Broadcom / AVGO	US / USD	321.68	-0.48%	2026-10-23 Stock Analysis • AVGO
TSMC ADR / TSM	US / USD	418.51	+1.49%	2026-10-23 Stock Analysis • TSM
NVIDIA / NVDA	US / USD	210.96	-0.59%	2026-10-23 Stock Analysis • NVDA
AMD / AMD	US / USD	516.79	+1.74%	2026-10-23 Stock Analysis • AMD
Intel / INTC	US / USD	86.96	+1.72%	2026-10-23 Stock Analysis • INTC
Cambricon / 688256	CN / CNY	932.20	+5.46%	2026-10-23 Stock Analysis • 688256
Hygon / 688041	CN / CNY	201.10	+5.28%	2026-10-23 Stock Analysis • 688041
SMIC / 00981	HK / HKD	47.06	+0.47%	2026-10-23 Stock Analysis • 00981

Networking, power and cooling

Company / ticker	Market / FX	Close	Change	Session / source
Arista Networks / ANET	US / USD	172.37	+0.76%	2026-10-23 Stock Analysis • ANET
Vertiv / VRT	US / USD	194.32	-0.65%	2026-10-23 Stock Analysis • VRT

Platforms and applications

Company / ticker	Market / FX	Close	Change	Session / source
Microsoft / MSFT	US / USD	524.51	+0.66%	2026-10-23 Stock Analysis • MSFT
Alphabet / GOOGL	US / USD	320.21	-1.24%	2026-10-23 Stock Analysis • GOOGL
Meta / META	US / USD	642.62	-0.96%	2026-10-23 Stock Analysis • META
Amazon / AMZN	US / USD	248.00	-0.95%	2026-10-23 Stock Analysis • AMZN
iFlytek / 002230	CN / CNY	38.20	+4.09%	2026-10-23 Stock Analysis • 002230
Baidu / 09888	HK / HKD	75.95	+1.13%	2026-10-23 Stock Analysis • 09888
Alibaba / 09988	HK / HKD	97.80	+0.46%	2026-10-23 Stock Analysis • 09988
Tencent / 00700	HK / HKD	386.00	+0.94%	2026-10-23 Stock Analysis • 00700
SenseTime / 00020	HK / HKD	1.16	+4.50%	2026-10-23 Stock Analysis • 00020

Next, watch for Kimi weights arriving on schedule, synthetic data improving real robot tasks, and capacity expansion and AI adoption producing verifiable operating results.