The Tech Times

MONDAY 2 NOVEMBER 2026 EVENING / 17:30 GMT

**THE TAKE**

**Useful AI needs more than a higher score.** A demanding new evaluation exposes uneven capabilities, while today's tool releases improve how agents reach data and run locally. The practical theme is controlled adoption: inspect the task, the migration path and the operational defaults.

Coverage: 2 November 2026, 05:30–17:30 UTC. Three fresh release groups; dated catch-ups and model-transition notices are labelled individually. Opening take: editorial interpretation.

Data and local AI / 2 · Build and automate / 3 · Upgrade carefully / 4 · API changes / 5

**EVALUATION CATCH-UP / PUBLISHED 2 NOV; TIME UNVERIFIED**

# GDPx tests agents on expert-domain work

- SemiAnalysis has introduced GDPx, with 206 tasks across 11 equally weighted domains with 14 specialist data partners. Its initial leaderboard puts Opus 5.5 at 12.0%, Gemini 4 Carbon at 10.9% and GPT-6 Astra at 9.6%. These are the benchmark author's results, not an independent reproduction.
- The aggregate mixes binary and continuous grading and averages three attempts. In finance and legal tasks, the intended binary grading gives zero even when partial-credit measures differ. Thresholds and domain weights therefore matter: 12% does not mean a model can perform 12% of all jobs, or is only 12% useful.
- For builders, the useful question is performance on your own domain, with your tools and cost limits - not the overall ranking alone. This is a dated catch-up absent from the morning edition; the primary announcement does not establish a post-05:30 publication time.

Announcement and methodology · Live leaderboard · Domain definitions

**UPCOMING / 4 & 6 NOV; NOTICE DATE UNVERIFIED**

# Together's next retirements arrive this week

- Together lists Qwen/Qwen3.8-2.4T-A95B for serverless removal on 4 November, recommending Qwen/Qwen3.7-Max. On 6 November, listed removals include MiniMaxAI/MiniMax-M3 and thinkingmachines/Inkling; suggested replacements are MiniMax-M3.1-Flash and Muse-Spark-Open-816B-A42B respectively.
- These are scheduled changes, not observed shutdowns or newly dated announcements. A recommended substitute does not establish equal cost or behaviour. Check hard-coded IDs and run a small compatibility test before changing production traffic.

Provider deprecation schedule

DATA AND LOCAL AI / FRESH RELEASES

DATA / RELEASED 2 NOV, 06:06 UTC

# Chartbrew adds an MCP route into dashboards

- Chartbrew 6.0 adds an MCP server for compatible AI tools to find permitted resources, explore data and prepare chart previews. It can also use external MCP services as data sources with OAuth. Its versioned Data API adds scoped keys and project controls; private charts can be exported as PNG without creating public links.
- The new Chart Studio keeps AI conversations with charts, and history retains up to 100 versions. But upgrading is not transparent: ECharts replaces Chart.js, the legacy chartData response field is removed, and the visual SQL builder is gone. Existing integrations must move to the new data contract or Data API. Source is available for self-hosting; feature-specific hosting and AI costs were not verified.

Version 6 release and migration changes · Setup and licence files

LOCAL INFERENCE / PRERELEASES / 2 NOV, 08:02-16:54 UTC

# llama.cpp tunes three GPU paths

- Build b12034 fixes WebGPU scratch sizing that could abort small-query-batch workloads. b12048 improves Intel Arc Pro B70 SYCL communication, while b12050 changes small-n Vulkan split-K selection for AMD RDNA3 and RDNA4.
- The Intel contributor reports 19% higher prompt-evaluation throughput in one tensor-split test with unchanged perplexity; Windows is excluded pending validation. This is not a general speed claim or our test. All three builds are public prereleases: compare the relevant hardware path with your current build before upgrading production.

WebGPU fix · Intel release · Intel test and limits · AMD release

AUTOMATION / PRERELEASE / 2 NOV, 08:27 UTC

# n8n repairs links through long pagination loops

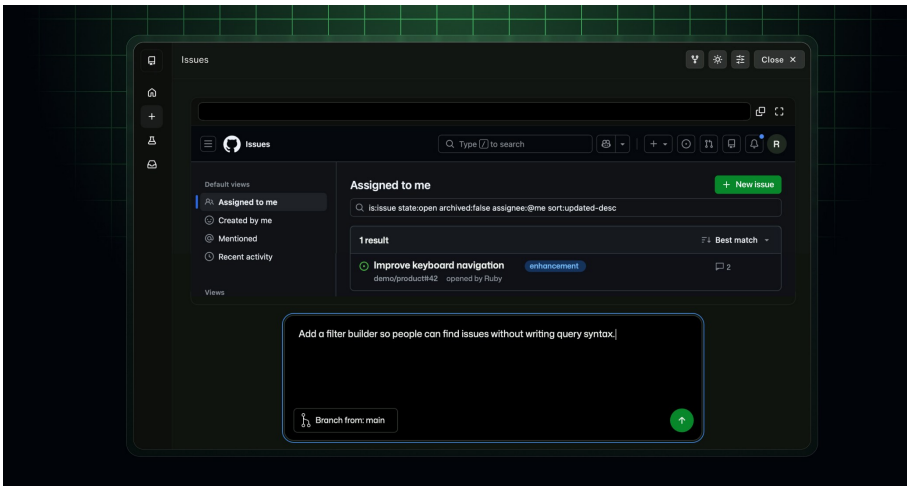
- n8n 2.46.4 fixes linked-item resolution across long pagination loops and preserves Error names and fields across the expression runtime. Migration-report scans also reuse user and Slack-credential reads and switch to keyset pagination. These are reliability and performance fixes, not a new AI feature. GitHub marks this build as a prerelease; test affected workflows before relying on it in production.

Official prerelease notes

**BUILD AND AUTOMATE / DATED CATCH-UPS**

**PRODUCT TOOLS CATCH-UP / ANNOUNCED 28 OCTOBER**

# GitHub opens a design-system-first prototyping tool




GitHub's Replica demonstration pairs a design request with a prototype. Official product image, not independently tested output. GitHub

- Replica turns a team's components, patterns and documentation into guidance for interactive prototypes in the GitHub Copilot app. The MIT-licensed tooling can run locally without a Hub; an optional self-hosted Hub provides publishing and comments. This is for exploring product ideas, not evidence that generated apps are production-ready.
- The documented local path needs Git, Node/npm and macOS or Linux. Model use needs a Copilot plan or a bring-your-own provider route that is still in public preview and carries provider charges. A self-hosted Hub does not mean all data stays inside the organisation: GitHub services and the selected AI provider have their own data handling.

Dated announcement · Requirements, licensing and privacy limits

---

**AGENTS CATCH-UP / GENERAL AVAILABILITY 28 OCTOBER**

# GitHub makes Markdown-defined agent workflows generally available

- GitHub Agentic Workflows compiles natural-language Markdown automation into Actions YAML for tasks such as issue triage, CI diagnosis and tests. General availability adds non-experimental spending forecasts, provider selection per workflow and self-hosted ARC/Docker-in-Docker support.
- It is available on GitHub Free, Team and Enterprise. The default Copilot CLI engine needs a Copilot licence; compute and model usage are not made free by the open-source tooling. Agents start read-only behind a sandbox/firewall, with validated outputs and a threat-detection job before proposed changes are applied. These safeguards are not a guarantee against all unsafe actions.

Official GA announcement and requirements

**UPGRADE CAREFULLY / DATED CATCH-UPS**

**SECURITY CATCH-UP / RELEASED 1 NOVEMBER**

# Gitea closes account-linking and Actions security gaps

- Gitea 28.3.0 fixes token-scope failures, OAuth account linking and Actions execution paths. Maintainers recommend upgrading promptly. In a non-default auto-linking configuration with open registration, an external identity could be linked to an existing account without that account's password; the advisory lists 28.3.0 as patched.
- A separate review-triggered workflow flaw exposed repository secrets and a write token to attacker-controlled workflow code. The release also repairs public-only token restrictions across several routes. These are disclosed vulnerabilities and patches - not evidence of exploitation in the wild. Back up and follow the project's upgrade process.

Release and fixes · Account-linking advisory · Actions advisory

**DATA CATCH-UP / RELEASED 1 NOV, 18:15 UTC**

## CocoIndex fixes keys that could overwrite the wrong data

- CocoIndex 1.1.0 audits connector identity rules after cases where distinct objects shared a key or an object changed identity between runs. Fixes cover same-named Drive files, cloud-storage path collisions, cross-topic Kafka keys and several database targets. The practical benefit is fewer silent overwrites, missed deletes and lost writes in indexing pipelines.
- Read the upgrade sequence before switching existing Neo4j or FalkorDB targets: the project warns that skipping the first of two updates deletes their nodes or edges. Multi-topic Kafka migration also changes keys and needs replay steps. Cloud-storage prefix matching now means a folder, so previously included similarly prefixed objects can be removed on the next update.

Release, affected connectors and migration warnings

**BROWSER AGENTS CATCH-UP / 1 NOV, 16:44 UTC**

## BrowserSkill reuses context without ignoring human intervention

- BrowserSkill 0.4.1 makes target-based shared control the default: DOM actions can reuse control context without a full-page resync after unrelated updates. Human intervention or navigation still revokes control, and visual actions and scripts retain stricter observation checks. This is not permission to continue after a person takes over.
- The CLI, running daemon, extension and DSH plugin must be updated together for protocol 1.8. Store review may delay the extension; a connection alone does not establish compatible versions. The release also removes automatic tab-group creation and its permission, without deleting existing groups.

Official release and coordinated upgrade steps

**API CHANGES / EFFECTIVE DATES AND UPCOMING DEADLINES**

**CREATOR API / EFFECTIVE 2 NOV; NOTICE 2 SEPT**

# An old image-model name now has a new default

- xAI's scheduled retirement of grok-imagine-image-quality reaches its effective date today. The official guide says the old name continues resolving to grok-imagine-image-2.0 at low quality for generation and editing; the already redirected pro name follows it. The original 1.0 model is unaffected.
- The documented low-quality rate is $0.01 less per image at every resolution than the retired model. Select the new model and quality explicitly rather than relying on an alias: omitting quality currently means low for generation but medium for editing. The response model field identifies what served the request. This reports the provider's migration policy, not a live API test.

Migration, defaults and pricing impact

**TRAINING API / EFFECTIVE 2 NOV; NOTICE 18 OCT**

# Fireworks changes its training pool

- Fireworks scheduled DeepSeek V4 Flash 0731 to leave its shared Serverless Training pool on 2 November. Its recommended replacement is DeepSeek V4.1 Flash 1017, offered there for LoRA training with up to 262K context. This notice is specifically about Serverless Training - not a shutdown of all inference, Managed Training or Dedicated Training. Check training configurations before the next job; actual removal was not tested.

Official dated training notice

**PUBLIC COVERAGE**

Public/indexed X, release notes, repositories and reporting were searched. This is not personal-feed access or exhaustive X coverage. No inaccessible post or engagement count is used as evidence. Original source links sit beside each story.

**EVIDENCE AND DATES**

Author-run benchmarks are not independent replications. Dated catch-ups retain their actual publication dates; effective-date notices are not new launches. Availability is documented, not hands-on tested. No unverified model launch is inferred from a leaderboard.