THE CURRENT

SOFTWARE / SCIENCE / MAKING AND MEASURING

Git learns to share the load

OpenAI’s Pando explains what changes when a code repository becomes infrastructure for a crowd of agents.

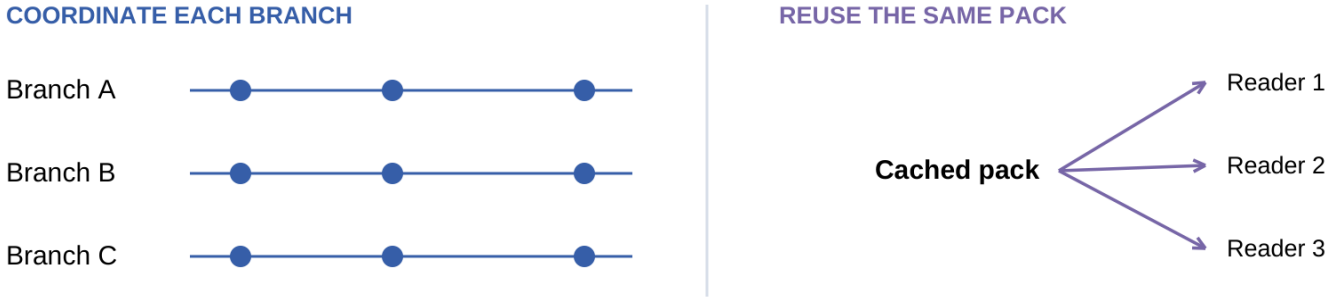

Original schematic of two Pando design ideas. Branches coordinate independently; readers reuse packed objects. It is not a latency benchmark.

A new report on a production system

On November 2, OpenAI described Pando, its internal Git backend. The company reports nearly fivefold better median push latency and almost a petabyte of daily Git traffic. This is an engineering disclosure about production infrastructure, not the launch of a Git hosting service. [1]

Two choices carry the story. Coordination moves to individual branches, so unrelated changes need not line up behind one repository-wide bottleneck. Fetches reuse cached packfiles instead of rebuilding a bespoke bundle for every reader. Some objects may be transferred unnecessarily, but that can be cheaper than repeatedly constructing the perfect response. [1]

The trade is explicit

Per-reference ordering remains, but readers do not receive one synchronized, point-in-time view across every reference. OpenAI’s interactive illustrations use synthetic workloads to explain the design; they are not production benchmarks. Public client work is available, while the backend is not a generally available service. [1, 2]

[1] OpenAI engineering report · Nov 2 | [2] OpenAI Git repository

Why this belongs beyond the Git desk

Editorial analysis: an agent fleet changes the shape of ordinary infrastructure demand. A human may pause between cloning, branching and pushing. Automated workers can create waves of similar operations. Under that workload, coordination and repeated preparation become as important as raw storage speed.

Pando offers a useful question for any shared system: which operations truly need to agree globally, and which merely happen to share a lock? Answering that question requires spelling out the consistency promise before removing the bottleneck. Speed without that contract can quietly change what users believe they have read.

In this edition

The same discipline runs through today’s issue: model findings need retesting; disc art depends on a real burning workflow; and an optical-switch demonstration is a promising component, not a finished machine. The interesting story is what each result actually makes possible.

RESEARCH METHODS

Put an expiry check on your assumptions

A study of aging LLM findings and a testing practitioner’s experiments point toward the same habit: measure again.



22 expired	17 weakened	8 persistent	3 strengthened
44%	34%	16%	6%

Fifty selected findings, classified in the study as of September 30, 2026. Each dot is one finding. “Expired” means at most 20% of the original effect, not zero. [3]

What survived newer models

A November 2 study by Stanford’s Owen Queen and James Zou revisits 50 findings across 20 models. The authors use minimal tests, sometimes substitute for unavailable original models, and fit trends across model releases. This is a selected sample, not a census of AI research. [3]

The tempting headline number needs its denominator. The reported 14.7-month median time to expiration applies only to the 22 expired findings. Across all 50, estimated expiration is 16% after one year and 39% after two. “Expired” also does not mean the behavior vanished in every model. [3]

Some problems strengthened, including judges’ preference for their own outputs and praise shaped by user preferences. Progress does not erase every weakness in parallel. [3]

An editorial maintenance rule

Treat a workaround as a claim with a date. Record the model, task, measurement and failure that justified it. When the model changes, rerun that small test before preserving the workaround indefinitely.

The testing recipe can age too

Dan Luu’s November analysis reports experiments mostly performed in September. Several structured testing approaches beat the default in the harder Pandoc experiment, but results depend on the task and prompt. The harder Pandoc experiment used 25 runs per condition; no condition achieved perfect holdout performance. [4]

Differential testing can make an agent compare independent implementations. Property testing can push it to identify invariants rather than memorize examples. Luu’s results suggest these techniques became more useful with the stronger model tested, without establishing one universally best recipe. TDD’s weaker showing here is not a verdict on TDD everywhere. [4]

A small useful experiment

Editorial proposal: keep a fixed set of representative changes and hidden checks. Compare the default workflow with one explicit testing method, using the same model and budget. Count correct outcomes, regressions, time and cost. Save failures as well as wins.

The goal is not a permanent ranking. It is to discover whether a method earns its extra work on today’s model and the problems you actually face.

[3] The shelf life of LLM findings · Nov 2 | [4] Dan Luu on testing · November

THE MAKER DESK

A picture written into a disc

DiscImageStudio expands its offline tools to macOS and Linux. The canvas is the recording surface itself.

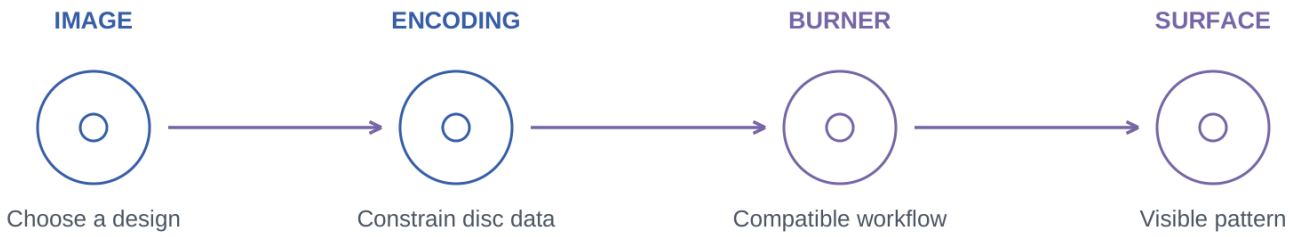

Original conceptual workflow. The software prepares an image; a compatible burning setup must still put the pattern onto physical media. No burn was performed for this issue.

The release and the older trick

Version 1.3.6, released November 2 at 13:52 UTC, adds experimental offline command-line packages for macOS and Linux on x64 and ARM64. The underlying disc-art technique predates this release. It encodes image patterns into data constrained by the disc's physical format, producing artwork visible on the recording surface after burning. [5, 6]

The project supports CD, DVD and single-layer 25 GB BD-R workflows. Its new non-Windows packages generate disc-image files; they should not be read as a promise of integrated live burning on every platform. A compatible burner and separate burning workflow remain part of the job. [5, 6]

A tool for experimentation

The Windows executable is unsigned, and live-stream burning is experimental. This is a maker project to approach with expendable media and careful compatibility checks, not a reason to risk an archive. Release availability has been checked; the binaries and physical results have not been tested here. [5, 6]

Why the physical trick is satisfying

Editorial reflection: software normally hides its physical substrate. We interact with filenames and pixels, not the arrangement of recorded regions beneath them. A visible disc pattern briefly reverses that relationship. The storage medium becomes part of the interface.

That makes the project more interesting than a novelty screenshot. Its constraints belong to the material: supported formats, writing behavior, hardware and the geometry of a surface. A successful file-generation step is only one part of a successful object.

A content seed worth keeping

"Where the bits become visible" could be a small illustrated series about the meeting point between information and matter. Start with optical-disc art, then examine one other medium through a real demonstration. Show the intermediate file, the physical process and the finished object separately.

The best version would preserve the failed attempts. Misaligned patterns and incompatible workflows teach more about the mechanism than a gallery of perfect examples alone.

[5] DiscImageStudio release · Nov 2 | [6] Project documentation and cautions

WORTH A CLOSER LOOK

Quantum engineering takes a networking detour

An October 28 optical-switch demonstration suggests a nearer-term use for difficult photonics work.

A component with its own destination

PsiQuantum has disclosed an 8-by-8 optical circuit-switch demonstration built on low-loss silicon photonics. The device redirects optical paths between input and output ports. Its laboratory measurements report typical insertion loss of 1.2 dB and reconfiguration below three milliseconds without amplification. These are company-reported results, not independent benchmarks. [7]

Selected prospective customers are receiving samples; commercial products are promised in future months. The announcement is neither a generally available switch launch nor evidence of a completed quantum computer. Its separate discussion of a 64-by-64 barium-titanate switch should not be conflated with this 8-by-8 demonstration. [7]

The concrete possibility is that components developed for quantum systems may also find uses in optical networking for AI infrastructure. Whether that becomes a competitive product depends on deployment requirements and evidence beyond this demonstration. [7]

The useful middle distance

Editorial reflection: a demanding research program can produce tools whose value does not depend on the entire original ambition arriving at once. That is a reason to examine components carefully, not to lower the standard of proof. A promising sample still has to become something a customer can integrate, operate and afford.

Gradient free is not feedback free

A November 3 critique by Matthew Honnibal adds useful precision to the weight-editing experiments published November 2. These explored whether an AI could change another model’s numerical weights directly to improve performance without gradient-based training. Different demonstrations used different setups. The 800-million-parameter Qwen experiment gave the editor a Python kernel, practice problems and answers, and a budget of forward passes. Gradients and optimizers were prohibited; feedback was not. [8, 9]

That distinction is documented in the original methods. It does not debunk the result, and the critique is not an independent rerun. It limits the story we should tell: a guided, gradient-free search is different from a model improving itself without tools or evaluation. [8, 9]

The last word

Editorial: availability is part of a technical result. Internal production infrastructure, downloadable experimental tools, selected-customer samples and a research protocol are four different kinds of opportunity. They deserve different next steps.

Read the system design. Try the maker tool cautiously. Watch the component’s independent tests. Reproduce the research on a bounded task. Curiosity gets stronger when it can name exactly what would change its mind.

[7] PsiQuantum optical switch · Oct 28 | [8] Honnibal critique · Nov 3 | [9] Weight editing experiment methods

Independent technology and science reading edition. Public sources only. Reporting checked November 3, 2026. Older material is dated explicitly; editorial proposals are not experiments performed.