THE CURRENT

SCIENCE / SOFTWARE / THE WORLD COMES INTO FOCUS

A hidden world comes into view

Rubin finds a distant dwarf planet candidate by connecting observations across years

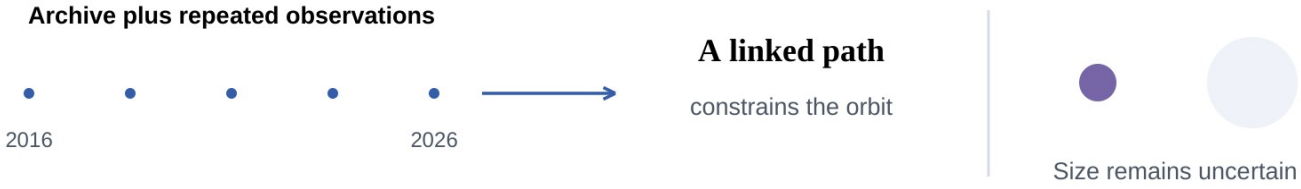

Original schematic. Linking positions constrains an orbit; brightness alone does not uniquely determine a world’s size.

Discovery announced October 27

A point of light lost in a crowded patch of sky has become a possible new member of the solar system’s small family of dwarf planets. Rubin Observatory announced 2026 ED138 on Tuesday: a distant body currently about 46 times farther from the Sun than Earth is, with an estimated diameter between 600 and 1,500 kilometers. The word candidate matters. Its size and classification still need follow-up. [1]

The discovery was a search through time as much as space. Matthew Holman identified the object on October 13 using software developed with Kevin Napier to connect moving detections. Thirty Rubin observations spanned twenty nights from March to July. Earlier DECam images extended the record back to 2016. What first looked like an unfamiliar dot now has more than a decade of orbital history. [1]

The region lies near the Milky Way’s crowded plane. Repeated observations let researchers separate an object’s motion from the dense background. This is the appeal of a survey that returns to the same sky: yesterday’s image can acquire a new meaning after tomorrow’s discovery. [2]

A bright dot is not a ruler

The diameter estimate depends partly on how reflective the surface is. A smaller, brighter surface and a larger, darker one can return similar amounts of sunlight. Infrared follow-up can help separate these possibilities. Until then, a midpoint such as 1,000 kilometers is an estimate, not a direct measurement. [2]

The scientific prize extends beyond naming another world. The sizes and orbits of trans-Neptunian objects help reconstruct how the outer solar system developed. The University of Washington’s account also links a public orbit viewer, making the geometry explorable while further measurements are gathered. [2]

Inside this issue

02 Apple gives agents a way into the document

A planned application interface may matter more than the new assistant.

03 A robot built not to ignite its surroundings

Hazardous-site inspections move toward customer delivery.

04 The terminal is a software boundary

Ghostty 1.4 brings graphics, speed and consequential fixes.

05 What kind of progress is this

A visual comparison and two short reflections.

[1] Rubin Observatory October 27 discovery report
[2] University of Washington Discovery and orbital context

Independent publication. No affiliation with or endorsement by organizations covered. Reporting cutoff October 28 at 15:00 UTC.

CREATIVE SOFTWARE

Apple gives agents a way into the document

November’s promised update makes structured access a more interesting story than another chat panel

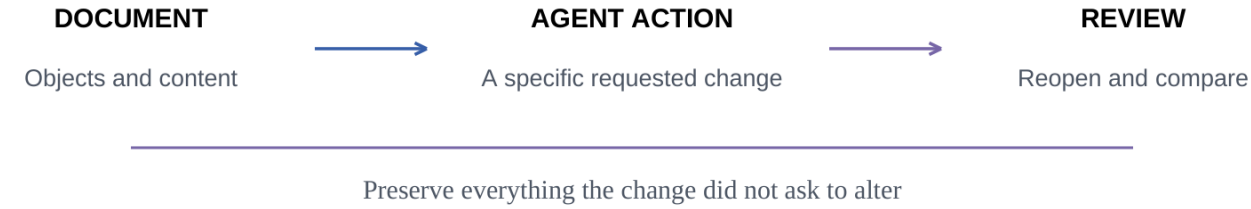

Conceptual workflow, not a confirmed implementation. The release announces application access; exact permissions and fidelity still need testing.

Announced October 27 Expected November

Apple’s next Creator Studio update will give external AI agents Model Context Protocol access to read, edit and perform actions in Pages, Numbers and Keynote. Keynote also gains .keyx, an open JSON-based bundle format intended for external tools and developer workflows. Those are planned capabilities, not features verified as available today. [3]

The suite also adds built-in assistants using Apple and third-party models, including OpenAI in the productivity apps. Apple explicitly describes Logic’s sound analysis as on-device; that does not establish that every assistant operation stays local. The assistants are beta and initially U.S. English, with supported hardware and OS requirements. MCP requires macOS 27 or later. [3]

Why an interface changes the job

An agent that manipulates a document through screenshots must infer what it sees, locate a control and hope the resulting state matches its intention. Structured access offers a different possibility: identify a particular object, request a change, then inspect the returned state.

That is an architectural opportunity, not a guarantee. The decisive question is whether the interface exposes meaningful document structure, rather than merely moving the same fragile clicking sequence behind a tool name.

The file is part of the interface

A readable interchange format could make presentation work easier to compare, generate and review outside its original app. A slide’s title, shape or chart might become something another tool can address directly. But JSON does not, by itself, promise faithful round trips, stable identifiers or preservation of every layout feature. Those are tests to run when the format and software ship.

A useful first trial would be deliberately ordinary: change a chart label, reorder two slides and replace one image in a copy of a real deck. Reopen it, compare the rendered pages and check speaker notes. The test should reward correct preservation, not merely the ability to produce a file that opens.

What creators can plan around

Apple says November updates are included for subscribers; new subscriptions cost $12.99 monthly or $129 annually in the U.S. Free versions of the productivity apps remain. [3]

For now, the sensible takeaway is to watch the technical documentation and release, not redesign an entire workflow around a press announcement. The missing evidence is concrete: permission controls, available operations, handling of unsupported objects and independent fidelity checks.

If those pieces hold together, the most valuable creative assistant may be the one that makes a small, exact edit while leaving the rest of the work alone.

[3] Apple Newsroom Creator Studio announcement and footnotes

ROBOTICS

A robot built not to ignite its surroundings

ANYmal X customer deliveries shift attention from impressive locomotion to the engineering of safe inspection



Conceptual protection layers described by the manufacturer. This diagram is not a certification guide or a claim of suitability for every hazardous area.

Customer delivery milestone October 27

ANYbotics says it has begun shipping ANYmal X to customers including Equinor and DSM-Firmenich. The quadruped is designed for inspection in explosive atmospheres, where a machine must do more than navigate a staircase: its own electronics, motors and battery must not become ignition sources. The announcement reports deliveries; it does not disclose unit counts or prices. [4]

The company describes nitrogen-filled chambers around computing and sensing equipment, flameproof housings for the twelve leg actuators and battery, conductive feet for grounding, and temperature cutoffs. These are engineering responses to specific failure paths, not a single magical property called “ex-proof.” [4]

Its product page lists Zone 1 IIB suitability and IP67 protection, with gas sensing optional. A buyer still needs to match the actual supplied configuration and certification to the site. Those labels should never be stretched into approval for every gas, dust or hazardous zone. [5]

Not the robot’s first walk

The earlier BASF Ludwigshafen case study describes twenty-minute missions with twenty inspection points, including gauge readings, thermal checks and pump acoustics. That 2023 field work matters: this week’s news is a commercial delivery step after early deployments, not the first demonstration that the machine can operate outside a laboratory. [6]

[4] ANYbotics October 27 customer delivery announcement
[5] ANYmal X Product scope and configuration
[6] BASF Ludwigshafen Earlier field deployment case study

The measurement is the product

ANYbotics connects inspection readings to asset records and maintenance workflows through its Shift platform. The broader system’s claimed 120,000 monthly inspections are not a count for the newly delivered X fleet. Independent confirmation of these shipments was not located for this edition. [4]

The practical value of a mobile inspector is repeatability. A thermal reading becomes more useful when it returns to the same pump, from a comparable viewpoint, under a recorded operating condition. A striking video of a walking robot says much less about whether a maintenance team can distinguish a real change from measurement noise.

That suggests better evaluation questions than “How autonomous is it?” How often does a route finish without intervention? How many readings are unusable? Can a technician see the evidence behind an alert? Does a missed inspection become visible, or simply disappear from the dashboard?

Watch the ordinary failures

Reliable inspection must accommodate the unglamorous cases: a blocked walkway, a dirty sensor window, a changed asset label or a battery that cannot finish the planned route. The relevant standard of success is a trustworthy inspection record over repeated operation, not one flawless demonstration.

DEVELOPER TOOLS

The terminal is a software boundary

Ghostty 1.4 improves the everyday experience while exposing why displayed output is not always passive

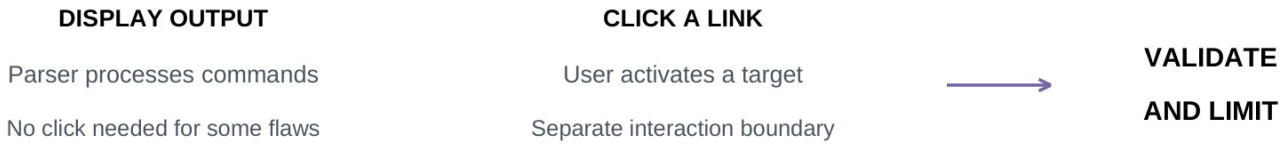

Different triggers carry different risks. The release fixes both output-processing flaws and separate interactions that require a click or paste.

Released October 27 Available now

Ghostty 1.4 is a substantial terminal release, with complete Kitty graphics support, including animation and relative image placement, plus memory and performance work. Programs can show animated images instead of a single frozen frame. Scrollback that is not visible can be compressed and restored when needed. [7]

The performance story has a useful mechanism behind it. On macOS, the maintainers separated draining the pseudo-terminal from parsing its output, reducing the cost of frequent small reads. Their M4 Max measurements show a 25–55% end-to-end throughput improvement over 1.3.x. These are maintainer benchmarks, not independently reproduced results or a promise that every command runs that much faster. [7]

The release is downloadable for macOS and through Linux packages or source routes. Existing users should also read the security notes: this is more than a cosmetic upgrade. [7, 8]

When output becomes input

One fixed vulnerability involved crafted Kitty graphics commands. The advisory describes crashes, excessive allocation and, in limited cases, unintended file reads or deletion. Displaying malicious output could trigger it; the user did not need to click a link. [9]

A separate macOS hyperlink flaw had a different trigger: Command-clicking a crafted link could open an unsafe local target or custom URL scheme and lead to command execution. Version 1.4 blocks unsafe targets and adds confirmation for custom schemes. Merely rendering that hyperlink was not the same execution path. [10]

Keeping those conditions separate matters. “A terminal vulnerability” can describe a parser problem, a clipboard action or a user-confirmed link. An accurate account tells the reader which boundary failed instead of flattening every issue into a frightening claim that any text executes code.

More useful also means more expressive

The practitioner discussion on Hacker News centers partly on this tradeoff: richer terminals make links and media convenient, while expanding the behavior their output can request. The discussion is useful criticism, not an independent security audit. [11]

A terminal is often treated as a plain-text window onto another program. In reality, its control sequences can request changes to state and presentation. The richer that vocabulary becomes, the more important validation, permission prompts and sensible defaults become.

For developers, the attractive part of 1.4 is that capability and hardening arrive together. For evaluators, the lesson is broader: test hostile output as well as friendly examples. A beautifully rendered result should not distract from the parser and policy that made it possible.

[7] Ghostty 1.4 Release notes and benchmark conditions
[8] Ghostty Official downloads
[9] Security advisory Kitty graphics output handling
[10] Security advisory macOS OSC 8 hyperlink handling
[11] Hacker News Practitioner discussion of Ghostty 1.4

COMPARISON AND ORIGINAL REFLECTION

What kind of progress is this

A discovery, a promise, a shipment and a release call for different next questions

DEVELOPMENT	CURRENT EVIDENCE	NEXT USEFUL CHECK
Rubin candidate	Observed and orbit linked	Constrain size and classification
Apple agent interfaces	Announced for November	Inspect the shipped interfaces
ANYmal X	Vendor reports deliveries	Verify site fit and route reliability
Ghostty 1.4	Released software and fixes	Check version and compatibility

The same word “new” can hide four different evidence states. Source details and qualifications appear in the preceding reports.

The value of looking again

A fresh discovery need not begin with a fresh observation. Sometimes the decisive move is to return to material that was already collected, now with a better question. An archive is not just a storehouse of finished knowledge. It is a reserve of possible relationships.

This is a useful way to think about personal learning, too. Saving more material is easy; creating a route back to the right material is harder. A note earns its place when a later question can find and use it.

The danger is to confuse accumulation with understanding. An enormous collection without a way to test connections can preserve error just as effectively as insight. The best archive is not the one that remembers everything equally. It is the one that makes a promising relationship inspectable.

That is the quiet pleasure in a discovery assembled across years: the world did not change to become interesting. Our ability to notice a pattern did.

Memory needs a maintenance phase

Research reflection June paper revisited

Google’s October 27 “Sleep” explainer describes alternating interaction with offline consolidation and generated practice. The associated paper was submitted in June and revised in July; this is an explanation of existing research, not a newly released model or API. [12, 13]

The idea is appealing because it separates two jobs: using information now and incorporating it for later. Its reported results are specific experiments, not proof of perfect general memory. In one mathematical evaluation, AIME 2025, the reported score rises from 68.1 to 69.0 for Qwen3-8B under the paper’s average-at-16 evaluation. That modest difference is a useful counterweight to sweeping interpretations. [13]

The larger design question is what should survive consolidation. A durable memory system needs to decide which examples deserve retention, which conflict with earlier evidence, and how a mistaken lesson can be corrected. More permanence is not automatically more intelligence.

A good memory is therefore partly a revision system. It should make room for learning without making every early impression irreversible.

[12] Google Research Sleep framework explainer October 27
[13] Language Models Need Sleep Paper and version history

Reporting draws on official releases, research and practitioner discussion. Vendor claims are attributed; this edition did not run product benchmarks or reproduce research results.