

← [rss-offline](#) · [abrir original](#) · **I think Anthropic and OpenAI have found product-market fit**
[rendered] · Simon Willison's Newsletter · 2026-05-28 23:19

I think Anthropic and OpenAI have found product-market fit

Plus Claude Opus 4.8: “a modest but tangible improvement”, and the Pope!



SIMON WILLISON

MAY 28, 2026

♡ 48

💬 5

🔄 3

Share

In this newsletter:

- I think Anthropic and OpenAI have found product-market fit
- Claude Opus 4.8: “a modest but tangible improvement”
- Notes on Pope Leo XIV’s encyclical on AI

Plus 5 links and 4 quotations and 1 note and 6 releases and 1 tool

Sponsor message: Accelerate your business growth by publishing your apps and agents on Microsoft Marketplace. Meet your customers where they are and connect with over 6 million customers around the globe 24/7 who trust the power of Microsoft. [Discover Microsoft Marketplace.](#)

I think Anthropic and OpenAI have found product-market fit - 2026-05-27

Anthropic are [strongly rumored](#) to be about to have their first profitable quarter. Stories [are circulating](#) of companies surprised at how expensive their LLM bills are

becoming from usage by their staff. I think this is because OpenAI and Anthropic have both found product-market fit.

- Enterprise customers are now paying API prices
- I think they've found product-market fit
- And they're ramping up
- The AI-failure stories around this are pretty thin
- We also know the labs are spending a lot
- API revenue is becoming less important
- April is a new inflection point

Enterprise customers are now paying API prices

I currently subscribe to the \$100/month Max plan from Anthropic and the \$100/month Pro plan from OpenAI. If you are a heavy user of coding agents these plans are a fantastic deal. I just ran the [ccusage](#) tool on my laptop to get an estimate of how much I would have spent if I were to pay for API tokens in the past 30 days and got:

- \$1,199.79 for Anthropic Claude Code
- \$980.37 for OpenAI Codex

That's \$2,180.16 worth of tokens for \$200 - not bad at all! I'm a moderately heavy user of these tools, but I'm certainly not running agents every hour of the day and night.

I had assumed that companies making extensive use of agents were getting similar discounts. It turns out I *could not have been more wrong* about that.

I haven't been able to track down the exact date, but at some point in the last six months Anthropic switched their Enterprise plan (originally "[Claude seats include enough usage for a typical workday](#)" [back in August 2025](#)) to \$20/seat/month plus API pricing for usage. This story about the change [from The Information](#) is dated Apr 14, 2026, but cites an Anthropic spokesperson claiming that the pricing change occurred

in November 2025. Existing customers are finding out about the change as they renew their contracts.

OpenAI made a similar pricing change in April. The [Codex rate card \(Internet Archive copy\)](#) currently says:

Note: On April 2, 2026, we updated Codex pricing to align with API token usage, instead of per-message pricing. This change was applicable to new and existing Plus, Pro, ChatGPT Business and new ChatGPT Enterprise plans.

On April 23, 2026, we made this update for all existing ChatGPT Enterprise plans as well, inclusive of Edu, Health, Gov, and ChatGPT for Teachers.

It's a little harder to decode as they quote prices in "credits", but as far as I can tell those credit costs are an exact match for the API token costs listed for those models.

All of which is to say that as of April 2026 the "Enterprise" cost for both OpenAI Codex and Anthropic Claude Code/Cowork is the same as the listed API price.

GPT-5.5 (released April 23rd) is 2x the API price of GPT-5.4. Opus 4.7 (April 16th) is [around 1.4x](#) the price of Opus 4.6 when you take their new tokenizer into account.

So April saw both leading model companies release new frontier models with a higher API price, *and* both companies now have measures to lock their enterprise customers (who tend to sign year-long deals) at those API prices, not the previous extreme discounts.

I think they've found product-market fit

Why these sudden aggressive moves on pricing? Both Anthropic and OpenAI are planning to IPO, but I suspect there's a more important factor here: I think they've finally found product-market fit, with the coding/general-purpose agent products embodied by Claude Code/Cowork and Codex.

Tools like ChatGPT are wildly popular, but that wild popularity has been difficult to turn into revenue. In February [OpenAI boasted](#) more than 900 million weekly active users for ChatGPT, but only 50 million - 5.6% of that - were paying consumer subscribers.

Charging \$10-\$20/month per user is an OK business, but you'd need 1-2 billion subscribers sticking around for four years to cover [\\$1 trillion in infrastructure](#).

Companies spending \$200+/month/user will get you there a whole lot faster - and as noted above, as a power-user I'm at ~\$1,000/month in API costs per vendor already.

Coding agents really did change everything. These are tools which burn *vastly* more tokens, but are also quickly becoming daily drivers for the work carried out by extremely well-compensated professionals. Right now that's still mostly software engineers, but a coding agent is a tool that can automate anything you can do by typing commands into a computer... so they are clearly applicable to a much wider set of skilled knowledge workers.

As I've [discussed on this site at length](#), the models released in November 2025 elevated agents to being genuinely useful. We've had six months to get used to that idea now - it's no wonder companies are beginning to spend real money on this technology.

You could argue that ChatGPT achieved product-market fit when it became the [fastest-growing consumer app in history](#) back in February 2023... but it certainly wasn't making any actual money back then. Coding agents plus enterprise pricing marks the point when these companies start making *very* real revenue. Maybe even enough to start covering their costs!

And they're ramping up

As further evidence that enterprise agents represent product-market fit for these companies, consider their open job listings.

OpenAI have [703 open jobs](#) right now, of which I'd categorize 229 (32.6%) as relating to enterprise sales and support - account executives, "Go To Market", "Forward Deployed Engineers" and the like.

Anthropic have [390 open jobs](#), 105 (26.9%) of which look enterprisey to me.

It's pleasingly ironic that these AI labs have picked a business model with such a heavy demand on human labor - enterprise sales contracts don't close themselves without a whole lot of humans in the mix!

(I ran this analysis by scraping their job sites with Claude Code, then having it use Datasette's [JSON API](#) to pipe that data into Datasette Cloud where I used [Datasette Agent](#) for the analysis, [exported here](#). Dogfood!)

The AI-failure stories around this are pretty thin

I started digging into this in response to [a growing volume](#) of stories claiming that large companies were sounding the alarm because their AI usage costs had grown so large.

The most widely cited of these stories appear quite overblown to me.

The most discussed has been Uber, based on [this report](#) where CTO Praveen Neppalli Naga indicated that Uber had "maxed out its full year AI budget just a few months into 2026", mostly thanks to Claude Code.

Given that Claude Code only got *really* good in November it's entirely unsurprising to me that a budget set in 2025 may have failed to predict demand for that tool in 2026!

That Uber story was further fueled by comments made by Uber's COO, Andrew Macdonald, on the Rapid Response podcast. I tracked down [the segment](#) and there really isn't much there. Here's what Andrew said:

But then you sometimes go and talk to your senior engineering leaders and you're saying, OK, how many projects that were on the cutting room floor got moved

above the line because of the productivity gains because 25% of our code commits were via Claude Code last quarter?

That link is not there yet, right? I think maybe implicitly there's more that is getting shipped. But it's very hard to draw a line between one of those stats and, OK, now we're actually producing like 25% more useful consumer features, right? And that line is hard to draw.

Somehow this fragment turned into headlines like [Uber's COO says it's getting harder to justify the money spent on AI tokenmaxxing](#), because the market for stories about AI failures remains enormous.

The other popular story around this is [Microsoft starts canceling Claude Code licenses](#), ostensibly to encourage their engineers to dogfood their own Copilot CLI agent instead - but The Verge reporter Tom Warren says "sources tell me the decision is also a financial one", triggered by the June 30th end of Microsoft's financial year.

I think both of these stories support my "product-market fit" hypothesis. The best advice I ever heard on pricing a product was that your customer should *suck air through their teeth* and then say yes. Uber's budget overrun and Microsoft's seat cancellations look like that effect playing out in practice.

We also know the labs are spending a lot

The big AI labs spend billions of dollars on both training and inference. Credible figures are hard to come by, but we did get one huge hint as to the figures involved from, oddly enough, the recent [SpaceX S-1](#):

[...] in May 2026, we entered into **Cloud Services Agreements with Anthropic PBC** ("Anthropic"), an AI research and development public benefit corporation, with respect to access to **compute capacity across COLOSSUS and COLOSSUS II**. Pursuant to these agreements, the customer **has agreed to pay us \$1.25 billion per month** through May 2029 [...]

The [Anthropic announcement](#) said that this deal meant they could “increase our usage limits for Claude Code and the Claude API”, heavily implying that Colossus is being used for inference, not model training.

Anthropic already have vast amounts of compute from other providers. The fact that they’re willing to spend \$1.25 billion per month for extra capacity from just *one* of their vendors hints at how big these inference budgets have become.

API revenue is becoming less important

Over the past two years my impression has been that OpenAI made more of their income from subscription revenue while Anthropic made more from their API.

Anthropic’s API revenue was historically quite dependent on a small number of large API customers - [this VentureBeat story from August 2025](#) quotes “sources familiar with the matter” suggesting that just Cursor and GitHub Copilot were responsible for \$1.2 billion of the company’s then-\$4 billion revenue.

Today Anthropic are rumored to hit [\\$10.9 billion in the second quarter](#), potentially even operating at a profit for the first time.

This pivot-to-Enterprise suggests that the labs have realized that the real money lies in cutting out the middlemen. Anthropic’s Claude Code directly competes with Cursor and Copilot. No wonder Cursor are [investing in their own models!](#)

April is a new inflection point

I’ve called November 2025 the [November inflection point](#) because that was when GPT-5.1 and Opus 4.5, combined with their respective coding agent harnesses, got *good* - good enough that we’ve spent the last six months adapting to agent systems that can reliably get useful work done.

I think April 2026 is a new inflection point where the revenue implications of this have started to land, to the benefit of the frontier AI labs and with material impacts on the budgets of large companies.

We'll know for sure how real this moment is when the S-1 documents for the upcoming Anthropic and OpenAI IPOs give us some real, audited numbers to get our teeth into.

Claude Opus 4.8: “a modest but tangible improvement” - 2026-05-28

Anthropic shipped [Claude Opus 4.8](#) today. My favourite thing about it is this note in the release announcement:

Users will find Opus 4.8 to be a modest but tangible improvement on its predecessor. There's still more to be done: we're working on developing and releasing models that provide many of the same capabilities as Opus at a lower cost.

It's so refreshing to see an AI lab honestly describe a release as a minor incremental improvement over the previous model!

Honesty seems to be a theme. Here's my other favorite note from that announcement:

One of the most prominent improvements in Opus 4.8 is its *honesty*. We train all our models to be honest---for instance, to avoid making claims that they can't support. But a general problem with AI models is that they sometimes jump to conclusions, confidently claiming to have made progress in their work despite the evidence being thin. Early testers report that Opus 4.8 is more likely to flag uncertainties about its work and less likely to make unsupported claims. This is borne out in [our evaluations](#), which show that Opus 4.8 is around four times less likely than its predecessor to allow flaws in code it has written to pass unremarked.

That linked system card includes the following:

Claude Opus 4.8 had the lowest incorrect-rate of the six models on every benchmark—the most direct measure of factual hallucination. It achieved this

mainly by abstaining on questions about which it was uncertain rather than by answering more questions correctly.

Model characteristics

Not much has changed since 4.7.

It's priced the same as Opus 4.5/4.6/4.7 - \$5/million input and \$25 per million output. "Fast mode" is twice that price, which is a significant reduction from their previous models - fast mode on 4.6/4.7 remains at \$30/\$150. Note that [fast mode](#) is only available to organizations that are part of the research preview, "Contact your account manager to request access".

Both the reliable knowledge cutoff and the training data cutoff are January 2026, the same as for 4.7.

The context window is still 1,000,000 tokens, and the max output is 128,000 tokens.

The [What's new in Claude Opus 4.8](#) document has some of the more interesting details. These caught my eye:

Mid-conversation system messages. Claude Opus 4.8 accepts `role: "system"` messages immediately after a user turn in the `messages` array (subject to [placement rules](#)). This lets you append updated instructions later in a long-running conversation without restating the full system prompt, which preserves [prompt cache](#) hits on the earlier turns and reduces input cost on agentic loops.

See also [this update](#) to the Anthropic Python SDK. Being able to steer the system prompt mid-conversation sounds really powerful. I was worried this would be incompatible with the abstraction provided by my own [LLM library](#), which expects a single system prompt per conversation... but it turns out my recent [redesign](#) should handle that [just fine](#).

Lower prompt cache minimum. The minimum cacheable prompt length on Claude Opus 4.8 is 1,024 tokens, lower than on Claude Opus 4.7.

I checked and 4.7's minimum was [4,096](#).

And some pelicans


Here are [pelicans riding bicycles](#) for all five thinking levels, low, medium, high, xhigh, and max:



low



medium



high



xhigh



max

This time I ran them using the [LLM CLI](#), exported the logs to Markdown and then had Claude Opus 4.8 [build me](#) an HTML tool that could render that Markdown with the `svg` fenced code blocks displayed as SVGs on the page.

(I later had GPT-5.5 xhigh in Codex [update that code](#) to remove any XSS holes. I'm sure Claude could have done that if I'd asked, but GPT-5.5 is my code security blanket at the moment.)

The max one was clearly the best, but it did take 25 input, 17,167 output tokens for a total cost of [43 cents](#)!

Notes on Pope Leo XIV's encyclical on AI - 2026-05-25

Dropped this morning by the Vatican: [Magnifica Humanitas of His Holiness Pope Leo XIV on Safeguarding the Human Person in the Time of Artificial Intelligence](#). This is a *very interesting* document. It's some of the clearest writing I've seen on the ethics of integrating AI into modern society.

Pope Leo XIV chose the name Leo in honor of Pope Leo XIII, who is known for his 1891 [Rerum novarum](#) encyclical on "Rights and Duties of Capital and Labor".

[This story](#) on Vatican News further clarifies the significance of that decision:

Meeting with the College of Cardinals for their first formal encounter after his election, Pope Leo XIV explained part of the reason for the choice of his papal name. "There are different reasons for this," he said, before going on to explain that he chose the name Leo "mainly because Pope Leo XIII, in his historic encyclical [Rerum novarum](#) addressed the social question in the context of the first great industrial revolution."

"In our own day," he continued, "the Church offers to everyone the treasury of her social teaching in response to another industrial revolution and to developments in

the field of artificial intelligence that pose new challenges for the defence of human dignity, justice, and labour.”

And now we get Pope Leo XIV’s own encyclical on the AI revolution. There’s a lot in here, but the writing style is very approachable, including to non-Catholics.

A few of my highlights

(I listened to most of the encyclical on a walk with our dog, my first time trying the [ElevenReader iPhone app](#). It worked very well: I pasted in a URL to the document and it read it to me in a very high quality voice, highlighting each paragraph as it went.)

Here are some of my highlights. In each case below **emphasis** is mine.

Here’s a useful description of the interpretability problem for LLMs in section 98:

First, any statement regarding AI risks becoming quickly outdated, given the remarkable pace at which these systems are developing. Second, all of us, including those who design them, possess only a limited understanding of their actual functioning. Indeed, **current AI systems are more “cultivated” than “built,” for developers do not directly design every detail, but instead create a framework within which the intelligence “grows.”** As a result, fundamental scientific aspects — such as the internal representations and computational processes of these systems — remain, at present, unknown.

I liked section 83’s description of the relationship between development and dignity:

For individuals as well as for nations, development is both a duty and a right. Minimum conditions are required for enabling every person and people to flourish in accord with their dignity, without being kept in a state of dependence or excluded from access to necessary goods. Development is truly human when it places people at the center instead of the accumulation of wealth, and when it concerns peoples as well as individuals. Justice demands the recognition of the rights of society and the rights of peoples, and includes a responsibility toward

future generations. **Development is not truly human if it increases consumption for some while shifting costs and burdens onto others, or relegates entire regions to subordinate roles, preventing them from realizing their full potential.**

Baked in cultural biases and sycophancy get a mention in section 100:

In personal use, three aspects in particular deserve careful consideration: the ease with which results are obtained, the impression of objectivity and the simulation of human communication. The speed and simplicity with which information, complex analyses, media content and practical assistance can be accessed undoubtedly makes life easier. Yet they can also encourage excessive reliance and the search for ready-made answers, and weaken personal creativity and judgment. **The apparent objectivity of the responses and suggestions these systems provide can lead us to overlook the fact that they reflect the cultural assumptions of those who designed and trained them, with all their strengths and limitations.** The artificial imitation of positive human communication — words of advice, empathy, friendship and even love — can be engaging and at times genuinely helpful. **However, for less discerning users, it can also be misleading, creating the illusion of a relationship with a real personal subject.** When words are simulated, they do not build genuine relationships, but only their appearance. The artificial imitation of care or support can become particularly risky when it enters contexts where real relationships and emotional bonds are lacking.

101 touches on the environmental impact:

Current AI systems require enormous amounts of energy and water, significantly influencing carbon dioxide emissions, and place heavy demands on natural resources. **As their complexity increases, especially in the case of large language models, the need for computing power and storage capacity grows too, which requires an extensive network of machines, cables, data centers and energy-intensive infrastructure.** For this reason, it is essential to develop more sustainable technological solutions that reduce environmental impact and help protect our common home.

102 covers the risks of algorithmic systems making decisions that impact people's lives without "compassion, mercy, forgiveness":

The use of AI is never a purely technical matter: **when it enters processes that affect people's lives, it touches on rights, opportunities, status and freedom.** Important and sensitive decisions — concerning employment, credit, access to public services or even a person's reputation — **risk being fully delegated to automated systems that do not know "compassion, mercy, forgiveness, and above all, the hope that people are able to change,"** and can therefore give rise to new forms of exclusion.

105 emphasizes the need for human accountability in how these systems are applied:

For AI to respect human dignity and truly serve the common good, responsibility must be clearly defined at every stage: **from those who design and develop these systems to those who use them and rely on them for concrete decisions.** In many cases, however, the internal processes leading to a result remain opaque, making it harder to assign responsibility and correct errors. **This is where accountability becomes crucial: the possibility of identifying who must "account" for decisions, justify them, monitor them, and, when necessary, challenge them and remedy any harm caused.**

And 108 touches on the way AI amplifies the power of those with resources:

In fact, as with every major technological shift, **AI tends to amplify the power of those who already possess economic resources, expertise and access to data.** In light of the common good and the universal destination of goods, this raises serious concerns, since small but highly influential groups can shape information and consumption patterns, influence democratic processes and steer economic dynamics to their own advantage, undermining social justice and solidarity among peoples. For this reason, it is essential that the use of AI, especially when it touches on public goods and fundamental rights, be guided by clear criteria and effective oversight, grounded in participation and subsidiarity.

That same section explicitly calls out data as something that should be thought of more as a public good:

[...] Moreover, **ownership of data cannot be left solely in private hands** but must be appropriately regulated. **Data is the product of many contributors and should not be treated as something to be sold off or entrusted to a select few.** It is necessary to think creatively in order to manage data as a common or shared good, in a spirit of participation, as [Saint John Paul II](#) already suggested regarding collective goods.

Given that Palantir is named after a *Lord of the Rings* reference, I can't help but wonder if the J.R.R. Tolkien quote from *The Return of the King* (section 213) was the Pope throwing a little shade at Peter Thiel.

The twentieth-century Catholic author J.R.R. Tolkien, in the words of a protagonist in one of his novels, described our responsibility in this way: “It is not our part to master all the tides of the world, but to do what is in us for the succour of those years wherein we are set, uprooting the evil in the fields that we know, so that those who live after may have clean earth to till.” The civilization of love will not arise from a single or spectacular gesture, but from the sum total of small and steadfast acts of fidelity that serve as a bulwark against dehumanization. For this reason, it is worthwhile pausing to reflect on some aspects of how we, each in our own way, can cooperate in building the civilization of love.

Another 2026 prediction down

On 6th January this year I joined the [Oxide and Friends 2026 predictions](#) podcast episode to talk about predictions for 2026, 2029 and 2032. I [wrote mine up here](#), with hindsight they weren't nearly ambitious enough - it's already undeniable that LLMs write good code, we've made huge advances in sandboxing and New Zealand kākāpō have indeed [had a truly excellent breeding season](#).

There's one segment from the episode that I didn't bother to include in my write-up, but that I can't resist providing as a lightly-edited transcript here:

Bryan Cantrill: [37:13](#)

I think that AI has created some real public perception problems for itself. And I think that you are gonna have one of the frontier model companies, this year, have a white paper explaining how the proliferation of AI will mean prosperity for everybody. They will be trying to make some economic argument - because this is gonna be a 2026 election issue, how we think of these things and how they are regulated and it's a big mess. There's more heat than light in this debate.

Simon Willison: [38:05](#)

I'd like to tag something on to that one: I think that only works if they can sort of wash that through existing trusted experts. Sam Altman and Dario are constantly publishing essays about this stuff and nobody believes a word they say. Get Barack Obama's signature on one of these position papers and *maybe* you've got something people might start to trust a little bit.

Adam Leventhal: [38:27](#)

Otherwise, it's just like "leaded gas is good for you", says Exxon.

Bryan Cantrill: [38:31](#)

I mean, yeah. God. Obama... let's go with that, that's a great one because if it's like Bill Clinton everyone's gonna kind of roll their eyes, so it's gotta be someone who's got real credibility saying that this is gonna be broad-based... I'd say if they get that person to do it, it's gonna be revealed that that's also a bit crooked.

Simon Willison: [38:57](#)

How about the Pope?

Bryan Cantrill: [39:01](#)

The Pope is very into this stuff! That's a great prediction. We've hit pay dirt. The Pope weighing in on LLMs and their economic impact on the world.

Simon, I'm giving you full credit if the Pope weighs in believing that this is gonna be economic devastation.

My prediction here looks a whole lot less insightful given the Leo XIV/Leo XIII relationship, which I was unaware of when we recorded the episode!

Release: [datasette-agent 0.1a3](#)

- “View SQL query” buttons for both visible tables and collapsed SQL result tool calls.
- Don't display empty reasoning chunks
- Improved handling of truncated responses - table still displays to the user even if the SQL results were truncated when showing the agent.

See [Datasette Agent, an extensible AI assistant for Datasette](#).

Release: [datasette-agent-charts 0.1a2](#)

- “View SQL query” buttons below rendered charts.

Link 2026-05-22 [The memory shortage is causing a repricing of consumer electronics:](#)

David Oks provides the clearest explanation I've seen yet of why consumer products that use memory are likely to get significantly more expensive over the next few years.

The short version is that memory manufacturers - of which there are just three remaining large companies - have a fixed capacity in terms of how many wafers they

can process at any one time. This fixed wafer capacity is then split between DDR - used in desktops and servers, LPDDR - used in mobile phones and low-energy devices, and HBM - used with GPUs.

Until recently, HBM got just 2% of that wafer allocation. The enormous growth in AI data centers has pushed that up to an expected 20% by the end of 2026, and “a single gigabyte of HBM consumes more than three times the wafer capacity that a gigabyte of DDR or LPDDR does”.

Memory companies have learned from the extinction of their rivals that you should always under-provision rather than over-provision your fabricator capacity. The profit margins and demand for HBM (high-bandwidth memory) will constrain the production of consumer-device RAM for several years.

This is already being felt in the sub-\$100 smartphone market, which is particularly important to markets like Africa and South Asia.

(The original title of the piece was “AI is killing the cheap smartphone” but I’m using the Hacker News rephrased title, which I think does more justice to the content.)

Link 2026-05-23 [On the <dl>](#):

I learned a few new-to-me things about the <dl> element from this article by Ben Meyer:

1. A <dt> can be followed by *multiple* <dd>
2. You can optionally group the <dt> and <dd> elements in a <div> for styling - but only a <div>.
3. You can label them using ARIA.
4. They’ve been called “description lists”, not “definition lists”, since [an HTML5 draft in 2008](#).

So this is valid:

```
<h2 id="credits">Credits</h2>
<dl aria-labelledby="credits">
  <div>
    <dt>Author</dt>
    <dd>Jeffrey Zeldman</dd>
    <dd>Ethan Marcotte</dd>
  </div>
</dl>
```

Here's a useful note from Adrian Roselli on [screen reader support for description lists](#).

Tool: [Mad House — Usborne Creepy Computer Games](#)

Via [Hacker News](#) I learned that UK publisher Usborne published [free PDFs of their 1980s Computer Books](#), some of which I remember working through on my Commodore 64 as a child.

These were so great! Beautifully illustrated books with fun projects made up of code you could type into your own machine.

I remember playing “Mad House” typed in from the 1983 book “Creepy Computer Games”, so I fed that PDF [into Claude](#) and had it build an interactive version of that game in JavaScript and HTML:

Build a vanilla JS artifact that exactly recreates the game Mad House from this book, make sure it's mobile friendly and has a suitable retro aesthetic

Credit the book title and link to <https://usborne.com/us/books/computer-and-coding-books>



Quote 2026-05-24

The most frustrating failure mode right now is that people submit issues that are not in their own voice. They contain an observed problem somewhere, but it has been thrown into a clanker and the clanker reworded it and made a huge mess of it. Typically, it was prompted so badly that the conclusions produced are more often than not inaccurate but always full of confidence. The result is complete guesswork on root causes, fake-minimal repros, suggested implementation strategies,

analogies to adjacent but often the wrong code, and long lists of error classes that might or might not matter. [...]

So at least personally, I increasingly want issue reports to be condensed to what the human actually observed:

1. I ran this command.
2. I expected this to happen.
3. This happened instead.
4. Here is the exact error or log.

[Armin Ronacher](#), on slop issues filed against [Pi](#)

Release: [datasette-fixtures 0.1a0](#)

One of the smaller features in [Datasette 1.0a30](#) is this:

New documented [datasette.fixtures.populate_fixture_database\(conn\)](#) helper for creating the fixture database tables used by Datasette's own tests, intended for plugin test suites.

This new plugin takes advantage of that API. You can try it out using `uvx` without even installing Datasette like this:


```
uvx --prerelease=allow \  
  --with datasette-fixtures datasette \  
  --get /fixtures/roadside_attractions.json
```

Which outputs:

```
{
  "ok": true,
  "next": null,
  "rows": [
    {"pk": 1, "name": "The Mystery Spot", "address": "465 Mystery Spot Road, Santa Cruz, CA 95065", "url": "https://www.mysteryspot.com/", "latitude": 37.0167, "longitude": -122.0024},
    {"pk": 2, "name": "Winchester Mystery House", "address": "525 South Winchester Boulevard, San Jose, CA 95128", "url": "https://winchestermysteryhouse.com/", "latitude": 37.3184, "longitude": -121.9511},
    {"pk": 3, "name": "Burlingame Museum of PEZ Memorabilia", "address": "214 California Drive, Burlingame, CA 94010", "url": null, "latitude": 37.5793, "longitude": -122.3442},
    {"pk": 4, "name": "Bigfoot Discovery Museum", "address": "5497 Highway 9, Felton, CA 95018", "url": "https://www.bigfootdiscoveryproject.com/", "latitude": 37.0414, "longitude": -122.0725}
  ],
  "truncated": false
}
```

Release: [datasette-agent 0.1a4](#)

Taking advantage of the new [makeJumpSections\(\)](#) JavaScript plugin hook added in [Datasette 1.0a30](#), datasette-agent now presents this “Start a new agent chat” interface as part of the Jump to menu, any time you hit /:

home simonw 

LIVE DEMO

Datasette Agent

An AI assistant for [Datasette](#) to help explore and analyze data in SQLite

[GitHub](#) [Discord](#) [Open Datasette Agent](#)

Privacy notice

This is a public demo. Anything you type into Datasette Agent, and any conversations or results it creates, may be reviewed and used by the site administrators as test data. Do not enter private, confidential, personal, or sensitive information.

ANNOUNCEMENT · 21ST MAY 2026

You can try this out by signing into agent.datasette.io using your GitHub account.

Release: [datasette 1.0a30](#)

The big new feature in this alpha is a new customizable “Jump to...” menu, described in detail in [The extensible “Jump to” menu in Datasette 1.0a30](#) on the Datasette blog. You can try it out by hitting / on latest.datasette.io - it looks like this:

home root ☰

Datasette Fixtures

[Homepage actions](#)

An example SQLite database demonstrating Datasette. [Sign in as root user](#)

Data license: [Apache License 2.0](#) · Data source: [tests/fixtures.py](#) · About: [About Datasette](#)

_memory
0 tables

counters
3 rows in 1 table
[counters](#)

ephemeral
0 tables

The new `jump_items_sql()` plugin hook allows plugins to add their own items to the set that's searched by the plugin.

Quote 2026-05-26

I cannot believe I'm saying this, but getting the literal Pope to canonize your product's specific technical limitations as a spiritual treatise is the single greatest act of vendor lobbying I have ever seen.

[Corey Quinn](#), on Anthropic co-founder Christopher Olah's [influence](#) on *Magnifica Humanitas*

Quote 2026-05-26

A lot of the emails I get from founders are now written in a hard-hitting journalistic style. I know they're written by AI, because no founder ever wrote this way before. And once you realize something is written by AI, it's hard not to ignore it.

I have never knowingly finished reading an email signed by a human but written by AI. It feels like being lied to, and who would stand for that?

[...] It makes me think less of the author. It means they can't write well unaided (or feel they can't), and that they're trying to trick me.

It's not impressive to use AI to write stuff for you; any teenager can do that.

[Paul Graham](#)

Link 2026-05-26 [Microsoft Copilot Cowork Exfiltrates Files](#):

The biggest challenge in designing agentic systems continues to be preventing them from enabling attackers to exfiltrate data.

In this case Microsoft Copilot Cowork (yes, that's [a real product name](#)) was allowing agents to send emails to the user's own inbox without approval... but those messages were then displayed in a way that could leak data to an attacker via rendered images:

Because these messages can contain external images that trigger network requests to external websites, data can be exfiltrated when a user opens a compromised message sent by the agent.

Since OneDrive can create pre-authenticated download links, a successful prompt injection could cause those links to be leaked, allowing files to be downloaded by the attacker.

Link 2026-05-26 [The pressure:](#)

Daniel Stenberg on the unprecedented level of pressure the curl team are facing right now thanks to the deluge of (credible) AI-assisted security issues being reported.

The rate of incoming security reports is 4-5 times higher than it was in 2024 and double the speed of 2025 -- meaning that **on average we now get more than one report per day**. The quality is way higher than ever before. The reports are typically *very* detailed and long. [...]

For the first time in my life, my wife voiced concerns about my work hours and my imbalanced work/life situation. I work more than I've done before, but the flood keeps coming. [...]

This is a never-before seen or experienced pressure on the curl project and its security team members. An avalanche of high priority work that trumps all other things in the project that is primarily mental because we certainly *could* ignore them all if we wanted, but we feel a responsibility, we have a conscience and we are proud about our work.

The good news is that curl is a very solid piece of software, so the vulnerabilities people are finding tend not to be of high severity:

What is also a good trend: almost no one finds *terrible* vulnerabilities. All vulnerabilities found the last few years in curl have *all* been deemed severity LOW or MEDIUM. I'm not saying there won't be any more HIGH ever, but at least they are rare. The [most recent severity high curl CVE](#) was published in October 2023.

Quote 2026-05-27

PICARD: Data, shields up

DATA: Brilliant! Shields can reduce damage we sustain. Not immunity. Not hubris. Just prudence. It's not precaution—it's strategy.

[camera shakes]

WOLF: HULL BREACHES ON NINE DECKS

DATA: Here's what happened: you told me to raise shields, and I didn't

[Kyle Ferrana](#), @KyleTrainEmoji

Link 2026-05-27 [sqlite AGENTS.md](#):

SQLite gained an AGENTS.md file [five days ago](#) - but it's not intended for their own development, it's presumably aimed at people who are pointing agents at the SQLite codebase. It includes:

SQLite does not accept pull requests without prior agreement and/or accompanying legal paperwork that places the pull request in the public domain. However, the human SQLite developers will review a concise and well-written pull request as a proof-of-concept prior to reimplementing the changes themselves.

SQLite does not accept agentic code. However the project will accept agentic bug reports that include a reproducible test case. Patches or pull requests demonstrating a possible fix, for documentation purposes, are welcomed.

The [most recent commit](#) to that file removed “(currently)” from “SQLite does not (currently) accept agentic code”, with the commit message “Strengthen the statement about not accepting agentic code”.

Meanwhile the SQLite forum was being flooded with so many AI-generated bug reports - of varying quality - that they've now [split those off](#) into a [new SQLite Bug](#)

[Forum](#). D. Richard Hipp is resolving issues on there with a flurry of commits to the codebase.

Release: [llm-anthropic 0.25.1](#)

- New model: [Claude Opus 4.8](#) (c laude - opus - 4 . 8).
- New -o fast 1 option for [fast mode](#), for organizations with that feature enabled on their account.
- Default max_tokens for each model now defaults to that model's maximum output rather than 8,192. [#72](#)

See also my [notes on Opus 4.8](#) - I used this new release of llm-anthropic to generate the pelicans.

Note [2026-05-29](#)

The most interesting thing about [Anthropic's \\$65B Series H announcement](#) is this line (emphasis mine):

Since our Series G in February, adoption has continued to grow across global enterprise customers, and our run-rate revenue crossed **\$47 billion** earlier this month.

Anthropic have made a bit of a habit of sharing their “run-rate revenue” in this kind of announcement, which is an annualized projection of their current revenue - typically calculated by taking the most recent month and multiplying by 12.

Earlier this year:

- Apr 6, 2026 in [Anthropic expands partnership with Google and Broadcom](#): “Our run-rate revenue has now surpassed **\$30 billion**—up from approximately **\$9**

billion at the end of 2025.”

- Feb 12, 2026 in [Anthropic raises \\$30 billion in Series G](#): “Today, our run-rate revenue is **\$14 billion**, with this figure growing over 10x annually in each of those past three years.”

I had [Claude Opus 4.8](#) make me this chart using [Matplotlib](#) (Claude: “a data line chart is more straightforward matplotlib work—not really a design piece”):



Back in April [Axios CEO Jim VandeHei](#) wrote that he could not find “any company — in any industry, in any era — that has scaled organic revenue this quickly at this level as Anthropic” - and that was when they were at a paltry \$30 billion.

(Also [in Axios today](#) is an anonymously sourced note that “An AI consultant tells Axios one of their clients recently spent half a billion dollars in a single month after failing to put usage limits on Claude licenses for employees” - times that by 12 and you get an extra \$6 billion in annualized run-rate!)

Ed Zitron was [extremely skeptical of that \\$30 billion number](#) - I wonder if his skepticism will update for the new \$47 billion figure.

I've seen a few people dismiss this as untrustworthy, because the numbers come from Anthropic. That doesn't hold up: these numbers were included in announcements of their fundraises, and lying to investors who just put in \$65 billion would be securities fraud. They're even less likely to lie given that the real numbers will no doubt come out in their S-1 when they file for their IPO.

If you find this newsletter useful, please consider [sponsoring me via GitHub](#). \$10/month and higher sponsors get a monthly newsletter with my summary of the most important trends of the past 30 days - here are previews from [January](#) and [February](#) and [March](#).

Subscribe to Simon Willison's Newsletter

Launched 3 years ago

AI, LLMs, web engineering, open source, data science, Datasette, SQLite, Python and more

Subscribe

By subscribing, you agree Substack's [Terms of Use](#), and acknowledge its [Information Collection Notice](#) and [Privacy Policy](#).



48 Likes • 3 Restacks

Discussion about this post

Comments

Restacks



Write a comment...



Atharva Wankhede 3h

just the title scares me. what do you mean - you can be a million dollar company without PMF? WTF? all the chase for PMF seems futile if you have big bro's backing.

♡ LIKE (1) 💬 REPLY

📄 SHARE



Mike 3h

I'd love to read Ed Zitron's <http://wheresyoured.at/> take on this <https://simonw.substack.com/i/199683963/enterprise-customers-are-now-paying-api-prices>


♡ LIKE 💬 REPLY

📄 SHARE


3 more comments...

[← rss-offline](#) · [abrir original](#) · **OpenCode Now Supports DigitalOcean Inference Router for Intelligent Model Routing [rendered]** · DigitalOcean Blog · 2026-05-28 18:02


Inside the Architecture of DigitalOcean's Inference Router




Introducing DigitalOcean AI-Native Cloud



[Now Available: Kimi K2.6](#)



[Missed the Deploy 2026 Keynote live? Watch it now and catch everything that was announced](#)



- [Blog](#)
- [Docs](#)
- [Careers](#)
- [Get Support](#)
- [Contact Sales](#)



- [Products](#) ▾

Featured AI Products



Compute

Build, deploy, and scale cloud compute resources



Containers and Images

Safely store and manage containers and backups

>

Managed Databases

Fully managed resources running popular database engines

>

Management and Dev Tools

Control infrastructure and gather insights

>

Networking

Secure and control traffic to apps

>

Security

Help protect your account and resources with these security features

>

Storage

Store and access any amount of data reliably in the cloud

>

[Browse all products](#)

- Solutions ▾

AI/ML

>

CMS

>

Data and IoT

>

Developer Tools

>

Gaming and Media

>

GPU

>

Hosting

>

Security and Networking

>

Startups and SMBs

>

Web and App Platforms

>

[See all solutions](#)

• Developers ▾

Community

>

Documentation

>

Developer Tools

>

Get Involved

>

Utilities and Help

>

• Partners ▾

Become a Partner

>

Marketplace

>

- [Pricing](#)
- [Log in](#)
- [Sign up](#)
- [Log in](#)
- [Sign up](#)



[Product updates](#)

OpenCode Now Supports DigitalOcean Inference Router for Intelligent Model Routing



By [Musa Malik](#)


AI/ML Engineer

- Updated: May 28, 2026
- 3 min read

[← Back to blog home](#)

Table of contents

1. [What is an Inference Router?](#)
2. [What Changed for OpenCode](#)
3. [Beyond Coding Agents](#)



[What is an Inference Router?What Changed for OpenCodeBeyond Coding Agents](#)

Coding agents today have a massive spending problem. Every request, whether you’re designing system architecture or writing a single-line docstring, often gets routed to the same expensive frontier model. The result: unnecessary token usage, higher inference costs, and little awareness of task complexity or budget constraints.

This high cost stems from a “one-size-fits-all” approach to model usage, where premium frontier models are utilized for trivial tasks that don’t require such intensive reasoning effort. In multi-agent workflows, where orchestrators delegate work to specialized subagents, this lack of discrimination frequently leads to runaway

costs and opaque failure modes. Without intelligent routing, developers can essentially be forced into closed-provider lock-in and high API usage fees, which quickly escalate during exploratory building phases.

[DigitalOcean Inference Router](#), now in Public Preview, was built to solve this problem by dynamically routing requests to the right model for the job. As part of DigitalOcean's AI-Native Cloud, it gives developers a unified way to control, optimize, and evaluate AI inference across models. And as of today, you can access it through [OpenCode](#), the open-source AI coding agent, in as little as a few seconds.

What is an Inference Router?

An Inference Router is the auto-mode pattern engineers are used to, but with deliberate control over the tradeoffs that matter: latency, cost, and output quality. Rather than statically pointing your coding agent to a single model, an Inference Router can analyze each request and route it to the model best suited for that specific task. Not the *most* powerful model available, but the *right* model. That distinction is what drives real savings without compromising on your desired quality of output.

To use DigitalOcean's Inference Router: Create an Inference Router from the [router catalog](#)—pick a preset or build a custom router via the API or UI. No GPU management, no infrastructure to run. Use it by setting “model”: “router:your-router-name” in any OpenAI-compatible API call.

What Changed for OpenCode

OpenCode has become one of the most popular AI coding harnesses on GitHub, [earning over 160,000 stars](#) by embracing a simple idea: developers should not be locked into a single model provider. Its rise has shown a demand for provider agnostic AI use cases. At Deploy 2026, Tyler Gillam - a core engineer on Inference Router - [demoed our integration live on stage](#), showing exactly how OpenCode and Inference Router work together to make intelligent model selection decisions in real time. If you want to see it before diving in, the full recording is linked at the bottom of this post.

Previously, integrating DigitalOcean models into OpenCode meant manually editing your `opencode.json`, adding each model by hand, a list that would be outdated within weeks given the pace of new model launches. So, we built a native OpenCode integration that supports Inference Routers and DigitalOcean Serverless Inference models right out of the box.

Now you can run the following steps:

1. Launch OpenCode (desktop, web or TUI) and run `/connect`
2. Select **Login with DigitalOcean**
3. Your Inference Routers are shown in the Model Selection tab

That's it. You're plugging directly into a routing layer that's already helping to make the cost & quality tradeoff decisions for you based on your stated needs — with our purpose-build Software Engineering preset.

Beyond Coding Agents

This integration is part of a broader effort to bring DigitalOcean's Inference Engine into the tools developers already use, while continuing to invest in open source and upstream contributions. OpenCode is one example of that direction.

The goal is to make intelligent, cost-aware model routing the default for coding agents, not something you have to manually configure and hope for the best. As the OSS model landscape keeps improving, routing intelligence

will become more valuable, not less. The gap between “frontier” and “good enough” is closing fast, and developers who take advantage of routing will consistently come ahead on both desired quality and cost.

If you’re using OpenCode, try `/connect` today. If you want to dig deeper on what Inference Router is and how it works, the full documentation is available below.

Inference Router Resources:

- [How We Built DigitalOcean Inference Router](#)
- [Inference Router Documentation](#)
- [OpenCode DigitalOcean Install](#)
- [InferenceRouter OpenCode Deploy 2026 Demo](#)

About the author



Musa Malik

Author

AI/ML Engineer

[See author profile](#)

Musa is an AI/ML Engineer on DigitalOcean's Agentic Inference Cloud team, working on Plano & Inference Router. He joined DigitalOcean in March 2026 through the acquisition of Katanemo Labs, where he was a core engineer. He writes about LLM infrastructure, agentic systems and developer experiences for AI applications.

[See author profile](#)

Share




-
- [Product Updates](#)

Start building today

From GPU-powered inference and Kubernetes to managed databases and storage, get everything you need to build, scale, and deploy intelligent applications.

[Sign up](#)



Related Articles




Product updates

Scalable, Cost-Efficient AI: Introducing Unified Batch Inference on DigitalOcean

[snameo](#)

- May 27, 2026
- 8 min read

[Read more](#)




Product updates

Request-Based Autoscaling Is Now Generally Available on App Platform

Bikram Gupta

- May 22, 2026
- 4 min read

[Read more](#)




Product updates

Powering the Inference Era: Inside the DigitalOcean AI-Native Cloud

Vinay Kumar, Chief Product & Technology Officer

- May 4, 2026
- 7 min read

[Read more](#)



Company

- [About](#)
- [Leadership](#)
- [Blog](#)
- [Careers](#)
- [Customers](#)
- [Partners](#)
- [Referral Program](#)
- [Affiliate Program](#)
- [Press](#)
- [Legal](#)
- [Privacy Policy](#)
- [Security](#)
- [Investor Relations](#)

Products

- [GPU Droplets](#)
- [Bare Metal GPUs](#)
- [Inference Engine](#)
- [Data & Learning](#)

- [Droplets](#)
- [Kubernetes](#)
- [Functions](#)
- [App Platform](#)
- [Load Balancers](#)
- [Managed Databases](#)
- [Spaces](#)
- [Block Storage](#)
- [Network File Storage](#)
- [API](#)
- [Uptime](#)
- [Cloud Security Posture Management \(CSPM\)](#)
- [Identity and Access Management \(IAM\)](#)
- [Cloudways](#)
- [View all Products](#)

Resources

- [Community Tutorials](#)
- [Community Q&A](#)
- [CSS-Tricks](#)
- [Write for DOnations](#)
- [Currents Research](#)
- [DigitalOcean Startups](#)
- [Wavemakers Program](#)
- [Compass Council](#)
- [Open Source](#)
- [Newsletter Signup](#)
- [Marketplace](#)
- [Pricing](#)
- [Pricing Calculator](#)
- [Documentation](#)
- [Release Notes](#)
- [Code of Conduct](#)
- [Shop Swag](#)

Solutions

- [AI GPU Hosting](#)
- [H100 Cloud GPU](#)
- [AI Training GPU](#)
- [GPU Inference](#)
- [VPS Hosting](#)
- [Website Hosting](#)
- [VPN](#)
- [Docker Hosting](#)
- [Node.js Hosting](#)
- [Web Mobile Apps](#)
- [WordPress Hosting](#)
- [Virtual Machines](#)
- [View all Solutions](#)

Contact

- [Support](#)
- [Sales](#)
- [Report Abuse](#)
- [System Status](#)
- [Share your ideas](#)



Company



Products



Resources



Solutions




Contact



 © 2026 DigitalOcean, LLC. [Sitemap](#).

- 
- 
- 
- 
- 
- 
- 
- 



 Dark mode is coming soon.

This site uses cookies and related technologies, as described in our [privacy policy](#), for purposes that may include site operation, analytics, enhanced user experience, or advertising. You may choose to consent to our use of these technologies, or manage your own preferences. Please visit our [cookie policy](#) for more information.

AGREE & PROCEED

DECLINE ALL

MANAGE CHOICES

[Cookie Preferences](#)

Home / News & insights / Company news

Still a developer. Just outside. Our latest GitHub Shop collection is here.

The ESC collection lets you escape the confines of your desk and get out into the sun where good ideas are bound to happen.



Lavinia Sfetcu · @laviniasf

May 28, 2026

 3 minutes

Share:   

/ Blog

That's what inspired us to create the new ESC collection. It's not a manifesto to put down tools and chill at the beach (though that sounds nice too), it's the recognition that occasionally we have to escape the confines of a desk for the problem-solving to begin.

[Check out the ESC collection](#)



Clothes. But make it developer.

Long-time fan of the shop already know our hat and socks. “But where’s the tee?” you all cried. You asked. [Here it is](#). Plus, the [hat](#) gets a fresh colorway. Plus, you can now walk around the pool *and* maintain that developer wink with our [new pool slides](#).

/ Blog



Get beach ready

Like to go loud? We have exactly the right outfit for you: our Cabana set of matching [shirt](#) and [shorts](#) with Mona, Copilot, and Rubber Ducky in full tropical mode. If you prefer something more understated, keep it simple with the [linen shirt](#). Either way, swing our [cooler tote](#) over your shoulder, and you're set.

/ Blog



The hoodie that keeps cool

The [black invertocat hoodie](#) is one of our bestsellers. Now you can get one... for your can. This [coozie](#) won't keep you warm, but it will keep your refreshing beverage of choice cool, in more ways than one.



/ Blog

Still developers. Just shopping

This is a collection for developers, by developers. So, of course, we had to have a little fun with the shopping experience too. We used a lidar scanner to create the background on the images. Head to the GitHub shop to play around with the backgrounds—change the colors, the zoom, and everything in between to fully personalize your shopping experience.

Consider this your `--yo!o` flag for the summer. The ESC collection is live at the GitHub Shop. Browse, add to cart, and take some time for yourself. Just you and a cold drink in a tiny hoodie.

Pssst... we've got something cool cooking for the World Cup too, check back soon.

Tags :

GitHub Shop

Written by



Lavinia Sfetcu

[@laviniasf](#)

More on GitHub Shop

The ultimate gift guide for the developer in your life

Finding the perfect gift for your favorite developer is easy with our top tips.


/ Blog

Code. Create. Commit. Welcome to dev/core

The new GitHub Shop collection is here. We're celebrating you.

Lavinia Sfetcu

Related posts



AI & ML

Take your local GitHub sessions anywhere

Kick off work in VS Code or the CLI, finish it from your phone. Remote control for GitHub Copilot sessions is now generally available on github.com and GitHub Mobile.

Candy Ho & Ellie Bennett

/ **Blog**



Company news

GitHub availability report: April 2026

In April, we experienced 10 incidents that resulted in degraded performance across GitHub services.

Jakub Oleksy



Company news

/ Blog

Joe Binder

Explore more from GitHub



Docs

Everything you need to master GitHub, all in one place.

Go to Docs [↗](#)



GitHub

Build what's next on GitHub, the place for anyone from anywhere to build anything.

Start building [↗](#)



/ Blog

Learn more [↗](#)



The GitHub Podcast

Catch up on the GitHub podcast, a show dedicated to the topics, trends, stories and culture in and around the open source developer community on GitHub.

Listen now [↗](#)

We do newsletters, too

Discover tips, technical guides, and best practices in our biweekly newsletter just for devs.

Your email address

* Your email address

Yes please, I'd like GitHub and affiliates to use my information for personalized communications, targeted advertising and campaign effectiveness. See the [GitHub Privacy Statement](#) for more details.

/ Blog

GitHub

Product

[Features](#)

[Security](#)

[Enterprise](#)

[Customer Stories](#)

[Pricing](#)

[Resources](#)

Platform

[Developer API](#)

[Partners](#)

[Atom](#)

[Electron](#)

[GitHub Desktop](#)

Support

[Docs](#)

[Community Forum](#)

[Training](#)

[Status](#)

[Contact](#)

Company

[About](#)

[Blog](#)

[Careers](#)

[Press](#)

[Shop](#)

© 2026 GitHub, Inc. [Terms](#) [Privacy](#) [Manage Cookies](#)
Do not share my personal information





[Stream](#) [Garden](#) [Now](#)  

May 28, 2026

How to add icons to external links with CSS

As part of my ongoing [updates to my personal site](#), I decided to add icons after all external links.

Today, I wanted to show you how I did it. Let's dig in!

Why?

I do a *lot* of linking on my site, and it's not always obvious when a link will keep you on-site vs. take you somewhere else.

Using `[target="_blank"]` forces links to open in a new tab or window, but this is considered an anti-pattern for accessibility reasons. It can be confusing, and it robs users of choice and agency.

An icon lets them make an informed decision.

Rather than manually adding a link to every external link, I figured out how to target all links that point to an external page and add an icon with CSS instead.

Targeting external links with CSS



Modern CSS makes this possible!

I want to target...

1. Links that start with `http` (*not* relative links that point to the current site), and
2. Do *not* have the `.btn` class (the class I use to style links as buttons), and
3. Do *not* have an SVG icon inside them already, and
4. Do *not* have my site URL or `localhost` in the URL.

I can use [the starts-with attribute selector \(`^=`\)](#) to check for links that start with specific patterns (like `http`).

The `:not()` and `:has()` pseudo-classes let me exclude certain selectors.

```
[href^="http"]:not(.btn, :has(svg)):not([href^="https://gomakethings.com"], [href^="http://localhost"]) {  
    /* ... */  
}
```

Adding Content

The `::after` pseudo-element lets me add content *after* the targeted link.

```
[href^="http"]:not(.btn, :has(svg)):not([href^="https://gomakethings.com"], [href^="http://localhost"]>::after {  
    /* ... */  
}
```

While you can add an SVG icon using `background-image`, you lose the ability to have it inherit fill color with the `currentColor` property and sizing relative to the target link.

To get around this, we need to instead use **the mask and background-color properties** [↗](#).

```
[href^="http"]:not(.btn, :has(svg)):not([href^="https://gomakethings.com"], [href^="http://localhost"]):after {
  background-color: currentColor;
  mask: url('data:image/svg+xml,<svg xmlns="http://www.w3.org/2000/svg" width="1em" height="1em" fill="currentColor" viewBox="0 0 16 16"><path fill-rule="evenodd" d="M8.636 3.5a.5.5 0 0 0-.5-.5H1.5A1.5 1.5 0 0 0 0 0 4.5v10A1.5 1.5 0 0 0 1.5 16h10a1.5 1.5 0 0 0 1.5-1.5V7.864a.5.5 0 0 0-1 0V14.5a.5.5 0 0 1-.5.5h-10a.5.5 0 0 1-.5-.5v-10a.5.5 0 0 1 .5-.5h6.636a.5.5 0 0 0 .5-.5"/><path fill-rule="evenodd" d="M16 .5a.5.5 0 0 0-.5-.5h-5a.5.5 0 0 0 0 1h3.793L6.146 9.146a.5.5 0 1 0 .708.708L15 1.707V5.5a.5.5 0 0 0 1 0z"/></svg>');
  mask-size: cover;
}
```

We also need to include an empty **content** string, a **display** property, and a **height** and **width** to ensure the icon actually occupies space in the DOM.

```
[href^="http"]:not(.btn, :has(svg)):not([href^="https://gomakethings.com"], [href^="http://localhost"]):after {
  background-color: currentColor;
  content: "";
  display: inline-block;
  mask: url('data:image/svg+xml,<svg xmlns="http://www.w3.org/2000/svg" width="1em" height="1em" fill="currentColor" viewBox="0 0 16 16"><path fill-rule="evenodd" d="M8.636 3.5a.5.5 0 0 0-.5-.5H1.5A1.5 1.5 0 0 0 0 0 4.5v10A1.5 1.5 0 0 0 1.5 16h10a1.5 1.5 0 0 0 1.5-1.5V7.864a.5.5 0 0 0-1 0V14.5a.5.5 0 0 1-.5.5h-10a.5.5 0 0 1-.5-.5v-10a.5.5 0 0 1 .5-.5h6.636a.5.5 0 0 0 .5-.5"/><path fill-rule="evenodd" d="M16 .5a.5.5 0 0 0-.5-.5h-5a.5.5 0 0 0 0 1h3.793L6.146 9.146a.5.5 0 1 0 .708.708L15 1.707V5.5a.5.5 0 0 0 1 0z"/></svg>');
  mask-size: cover;
}
```

```
height: 0.8em;
width: 0.8em;
}
```

Odds & Ends

A few final-but-important details.

The icon may visually indicate “external link” to sighted users, but will not announce anything to visually impaired users. The `content` property has an *alt text* option: `<content> / <alt text>`.

Let’s add that with `(external link)` as the text.

```
[href^="http"]:not(.btn, :has(svg)):not([href^="https://gomakethings.com"], [href^="http://localhost"]):after {
  background-color: currentColor;
  content: "" / "(external link)";
  display: inline-block;
  mask: url('data:image/svg+xml,<svg xmlns="http://www.w3.org/2000/svg" width="1em" height="1em" fill="currentColor" viewBox="0 0 16 16"><path fill-rule="evenodd" d="M8.636 3.5a.5.5 0 0 0-.5-.5H1.5A1.5 1.5 0 0 0 0 0 4.5v10A1.5 1.5 0 0 0 1.5 16h10a1.5 1.5 0 0 0 1.5-1.5V7.864a.5.5 0 0 0-1 0V14.5a.5.5 0 0 1-.5.5h-10a.5.5 0 0 1-.5-.5v-10a.5.5 0 0 1.5-.5h6.636a.5.5 0 0 .5-.5"/><path fill-rule="evenodd" d="M16 .5a.5.5 0 0 0-.5-.5h-5a.5.5 0 0 0 0 1h3.793L6.146 9.146a.5.5 0 1 0 .708.708L15 1.707V5.5a.5.5 0 0 0 1 0z"/></svg>');
  mask-size: cover;
  height: 0.8em;
  width: 0.8em;
}
```

The icon I chose sits slightly off center. I use some margins to nudge-and-tweak it to where I want it to actually be.

```
[href^="http"]:not(.btn, :has(svg)):not([href^="https://gomakethings.com"], [href^="http://localhost"]):after {
  background-color: currentColor;
  content: " / "(external link)";
  display: inline-block;
  margin-block-end: -0.1ch;
  margin-inline-start: 0.5ch;
  mask: url('data:image/svg+xml,<svg xmlns="http://www.w3.org/2000/svg" width="1em" height="1em" fill="currentColor" viewBox="0 0 16 16"><path fill-rule="evenodd" d="M8.636 3.5a.5.5 0 0 0-.5-.5H1.5A1.5 1.5 0 0 0 0 0 4.5v10A1.5 1.5 0 0 0 1.5 16h10a1.5 1.5 0 0 0 1.5-1.5V7.864a.5.5 0 0 0-1 0V14.5a.5.5 0 0 1-.5.5h-10a.5.5 0 0 1-.5-.5v-10a.5.5 0 0 1 .5-.5h6.636a.5.5 0 0 0 .5-.5"/><path fill-rule="evenodd" d="M16 .5a.5.5 0 0 0-.5-.5h-5a.5.5 0 0 0 0 1h3.793L6.146 9.146a.5.5 0 1 0 .708.708L15 1.707V5.5a.5.5 0 0 0 1 0z"/></svg>');
  mask-size: cover;
  height: 0.8em;
  width: 0.8em;
}
```


Here's a demo on CodePen that you can play with! [↗](#)

Feel free to steal this, swap in your own icon, adjust the target URLs, and so on.


 **Was this helpful?** A [Go Make Things Membership](#) is the best way to support my work, and helps keep my content free and open to everyone.

 [Subscribe with RSS](#)

Stream Garden Now   

Made with  in Massachusetts. Unless otherwise noted, all code is free to use under the [MIT License](#).

[← rss-offline](#) · [abrir original](#) · **Why Japanese companies do so many different things [rendered]** · Leadership in Tech · 2026-05-28 03:05



Leadership in Tech

Why Japanese companies do so many different things

#315 – May 28, 2026

how toilet company became a major supplier of a critical semiconductor component

If you are a human, ignore this field

[Why Japanese companies do so many different things](#)

26 minutes by David Oks

Toto, the Japanese toilet company, became a major supplier of a critical semiconductor component called the electrostatic chuck. Demand from AI data centers has driven this once minor division to become Toto's largest and most profitable business. This reflects something broader about Japanese companies: they tend to do many very different things at a high level of quality. Economists trace this to a distinct bundle of corporate practices built around lifetime employment, horizontal coordination, and reinvestment over shareholder returns, which excels at incremental refinement but struggles with radical innovation.

[Still writing tests manually? Meticulous AI is here.](#)

sponsored by Meticulous

Most teams are forced to make the tradeoff between better coverage and more maintenance effort. But top engineering teams like Notion, Dropbox, Wealthsimple and LaunchDarkly have discovered a new testing paradigm. Built by ex-Palantir engineers, Meticulous creates and maintains E2E UI tests that covers every edge case of your web app - without any developer effort, making it the only tool to improve both product quality and dev velocity.

[The ask](#)

5 minutes by Michael Lopp

Michael explains how experienced leaders handle unexpected meetings by identifying “The Ask” — the real reason the meeting exists. Through examples, he shows that leadership is not only about strategy and data, but also communication, trust, instinct, and long-term relationships. Effective leaders learn to uncover hidden needs, guide collaboration, and rely on experience-based intuition to make better decisions and build connections across teams.

[Cheap turpentine](#)

20 minutes by David Singleton

David shares the actual documents and rituals that ran his 18-person AI startup, not frameworks but real artifacts: offer letters, hiring processes, board memos, and weekly rhythms. The team kept meetings minimal, shipped code daily, and ran Friday hackathons where everyone logged product friction that fed directly into the next week's priorities. Transparency was radical, with full board decks and readouts shared openly with the whole team.

[Prioritization happens in layers](#)

7 minutes by Ant Murphy

Prioritization is not a single backlog exercise but a layered process that starts with vision, strategy, and outcomes before reaching ideas and tasks. Many teams struggle because they try to prioritize isolated backlog items without clear direction. By aligning decisions across these layers, teams can reduce complexity, focus on the highest-impact opportunities, and use frameworks like RICE more effectively within the right context.

[Engineering metrics for beginners](#)

10 minutes by Joseph Gefroh

Metrics are easy to misuse. The right approach is to track them only when you plan to act on them, and to understand what they actually measure before using them to judge performance. For engineering teams, three areas matter most: delivery pipeline health, system reliability, and product observability. Key delivery signals include deployment frequency, failure rates, and cycle time. For reliability, focus on uptime targets that fit your actual needs and fast incident recovery over avoiding incidents entirely.

And the most popular article from the last issue was:

- [Barely treading water](#)

If you are a human, ignore this field

newsletters

- [Programming Digest](#)
- [React Digest](#)
- [C# Digest](#)

© 2013-2026 [Bonobo Press](#)
[Newsletters](#) · [Privacy](#) · [Advertise](#)

Changelog

Updates to the Tailscale client and service.

[Subscribe via RSS](#)
[All Changes](#)[Client](#)[Service](#)[Other](#)

May 28, 2026

Tailscale v1.98.4

[Update instructions](#) →

All Platforms

- An issue causing a deadlock when processing peer changes and disconnecting from the Tailscale control server is resolved.

May 21, 2026

Tailscale v1.98.3

[Update instructions](#) →

Linux

- An issue where [netfilter](#) rules could be applied inconsistently after a netfilter mode change failed has been resolved. Connmark and CGNAT rules are applied after the active netfilter mode is successfully updated, matching the behavior of other netfilter paths.

May 15, 2026

Tailscale v1.98.2

[Update instructions](#) →

All Platforms

- Go is updated to 1.26.3.
- A regression from 1.98.0 is patched, ensuring MagicDNS resolves tailnet hostnames after a network change. Windows clients are unaffected.

May 13, 2026

Tailscale v1.98.1

[Update instructions](#) →

Note: 1.98.0 was a [release candidate](#) intended for testing only.

All Platforms

- Expired [preferred peer addresses](#) are cleared, reducing the time to elect an alternative.

Linux

Note: 1.98.1 introduced a regression in the interaction between Tailscale and MagicDNS on Linux. The Linux release has been withdrawn pending a fix.

- [IP forwarding](#) for [subnet routers](#) and [exit nodes](#) is reported as a health check, surfacing warnings when misconfigured.
- `src_valid_mark` sysctl is set in addition to connmark firewall rules to improve [reverse path filtering](#) and prevent dropped packets.

macOS

- Devices and exit nodes can be searched from the application's main menu on macOS Tahoe 26 or later.
- The `AppIntroShown` system policy disables the **Welcome to the Tailscale app** modal window introduction that appears when you log in to Tailscale on a device for the first time.
- The **Hide Dock Icon** checkbox in Settings is available when Tailscale is disconnected.
- The `tailscale drive` CLI command directs users to configure [Taildrive](#) through the client GUI instead of returning an error.
- Device list is more responsive, especially for larger [tailnets](#).

iOS

- iOS devices can be used as [exit nodes](#).
- Device list is more responsive, especially for larger [tailnets](#).

Mar 30, 2026

Tailscale v1.96.5

[Update instructions](#) →

iOS

- An issue that could cause the network extension to encounter an out of memory condition on large tailnets is resolved.

tvOS

- An issue that could cause the network extension to encounter an out of memory condition on large tailnets is resolved.

Mar 27, 2026

Tailscale v1.96.4

[Update instructions](#) →

Linux

- An issue on forks of Linux caused by fallback-on-ENOSYS logic is resolved.
- An issue that could cause a segmentation violation during startup on MIPS devices is resolved.

Android

- An issue causing a deadlock when disconnecting from a tailnet is resolved.

Synology

- An issue on forks of Synology Linux caused by fallback-on-ENOSYS logic is resolved.

Mar 19, 2026

Tailscale v1.96.3

[Update instructions](#) →

Windows

- [DNS](#) resolution issue caused by NRPT rule formatting is resolved.

Mar 18, 2026

Tailscale v1.96.2

[Update instructions](#) →

Note: 1.96.0 and 1.96.1 were [release candidates](#) intended for testing only.

All Platforms

- `tailscale dns query|status` command supports `--json` flag to return JSON output.
- `tailscale wait [flags]` command waits for Tailscale resources to become available for binding.
- `tailscale ip` command supports `--assert=<specific-ip-address>` flag to assert that one or more of the node's IP addresses matches the specified IP address.
- `tailscale version --track` and `tailscale update --track` support release-candidate flag to check for and update to [release candidate builds](#).
- For 1.96.x, Go is updated from version 1.25 to 1.26.
- [Tailscale Peer Relays](#) advertise addresses discovered via Amazon EC2 Instance Metadata Service.
- `tailscaled_peer_relay_endpoints` gauge [user metrics](#) are available for [Tailscale Peer Relays](#).
- The AuthKey [system policy](#) applies only when a user is not in a logged in state.
- [UPnP](#) routes as expected during long lived [port mapping sessions](#) scenarios, including [hard NAT](#).

Linux

- Launch the systray application on startup using autostart file with the `tailscale configure systray --enable-startup=freedesktop` command.
- Scaling of [Tailscale Peer Relays](#) UDP sockets is gated by container-aware GOMAXPROCS defaults.
- Firewall rules created on Linux platforms correctly mark their traffic, avoiding [reverse path filtering](#) dropping connections and producing health warnings and risk prompts.
- OpenWrt versions 25.12.0 or later using apk as a package manager supports Tailscale updates.

macOS

- Windowed UI mode for macOS is [generally available](#).
- Double-click an account in the **Accounts** section to switch to that account.
- A progress dialog indicates Tailscale is waiting on the browser to complete reauthentication.
- The [open source variant](#) of Tailscale on macOS sets the `node:osVersion` attribute.
- The [Taildrop Send File](#) action and shortcut do not transmit empty files on macOS Tahoe (version 26) or later.
- Tailscale data directories for the [macOS standalone version](#) are excluded from Time Machine backups.
- An issue that required a machine reboot after installing a Tailscale update is resolved.

iOS

- iOS [bug report](#) ID displays in its entirety instead of being truncated.
- The [Taildrop Send File](#) action and shortcut do not transmit empty files on iOS version 26 or later.

Feb 12, 2026

Tailscale v1.94.2

[Update instructions](#) →

All Platforms

- Memory leak caused by high network map response rates is resolved.

Jan 26, 2026

Tailscale v1.94.1

[Update instructions](#) →

Note: 1.94.0 was a [release candidate](#) intended for testing only.

All Platforms

- `tailscaled_home derp_region_id` [client metrics](#) are available.
- `tailscaled_peer_relay_forwarded_packets_total` and `tailscaled_peer_relay_forwarded_bytes_total` [client metrics](#) are available for [Tailscale Peer Relays](#).
- [Identity tokens](#) are automatically generated for [workload identities](#).
- `--audience` flag added to `tailscale up` command to support auto generation of ID tokens for workload identity.
- `tsnet` [nodes can host Tailscale Services](#).
- The `tailscale lock status -json` command returns [tailnet key authority \(TKA\)](#) data in a stable format.
- [Tailscale Peer Relays](#) deliver improved throughput through monotonic time comparison optimizations and reduced lock contention.
- [Tailscale Services](#) virtual IPs are now automatically accepted by clients across all platforms regardless of the status of the `--accept-routes` feature.

Linux

- [Custom DERP servers](#) support Google Cloud Platform (GCP) Certificate Manager.
- [Tailscale SSH](#) authentication, when successful, results in LOGIN audit messages being sent to the kernel audit subsystem.
- [Tailscale Peer Relay](#) throughput is improved when the SO_REUSEPORT socket option is supported on multi-core systems.
- [Tailscale Peer Relay](#) server handshake transmission is guarded against routing loops over Tailscale.
- MagicDNS always resolves when using `resolve.conf` without a DNS manager.

macOS

- [AuthBrowser.macos](#) system policy sets a preferred browser for opening automatic authentication URLs.
- [HideDockIcon](#) system policy determines if the Tailscale Dock icon persists after all Tailscale windows close.
- Install and automatically update to [release candidate](#) versions of the client in the **About** section, **Release Channel** drop-down.
- DNS related health warnings no longer display when [Tailscale DNS](#) is disabled.
- `tssentinelId` command injection vulnerability has been removed. This fix addresses a security vulnerability described in [TS-2026-001](#).
- Ping view is [Tailscale Peer Relay](#) aware.

iOS

- Ping view is [Tailscale Peer Relay](#) aware.

tvOS

- **Use Tailscale Subnets** toggle is added in **Subnet Routing Settings**.
- Ping view is [Tailscale Peer Relay](#) aware.

Android

- Ping view is [Tailscale Peer Relay](#) aware.

Jan 6, 2026

Tailscale v1.92.5

[Update instructions](#) →

Linux

- [State file encryption](#) and hardware attestation keys are no longer enabled by default.
- Failure to load hardware attestation keys no longer prevents the client from starting. This could happen when the TPM device is reset or replaced.

Windows

- [State file encryption](#) and hardware attestation keys are no longer enabled by default.
- Failure to load hardware attestation keys no longer prevents the client from starting. This could happen when the TPM device is reset or replaced.

Dec 16, 2025

Tailscale v1.92.3

[Update instructions](#) →

All platforms

- [WireGuard](#) configuration that occurs automatically in the client, no longer results in a panic.

macOS

- [Tailscale system extension](#) no longer fails to install during an upgrade.

Dec 10, 2025

Tailscale v1.92.2

[Update instructions](#) →

macOS

- [Tailandrop](#) works as expected using the macOS **Share** option.

Android

- An issue in [custom control servers](#) (Headscale) that could result in connectivity problems is resolved.

Tailscale v1.92.1

[Update instructions](#) →

Note: 1.92.0 was a [release candidate](#) intended for testing only.

All platforms

- [Tailscale Funnel](#) and [Tailscale Serve](#) support the PROXY protocol, a header format that forwards information about the original client connection, such as the source IP and port, to the server before the actual traffic begins.
- [Tailscale Peer Relays](#) can use static endpoints using the `tailscale set` command with the `--relay-server-static-endpoints` flag.
- [Tailscale Services](#) can be configured to use a remote target as a service destination.
- Nodes can authenticate using [workload identity federation](#) with the `tailscale up` command flags `--client-id` and `--id-token`.
- [Network flow logs](#) automatically record node information about itself and peers it communicates with.
- [Tailnet Lock](#) command `tailscale lock log --json` response returns [Authority Update Messages](#) (AUMs) in a more stable format.
- Tailscale Peer Relay endpoint advertisements include more candidate IP:port pairs.

- Tailscale Peer Relays support multiple, forward bind packets per handshake generation, which improves path selection and chances of completing a handshake.

macOS

- Redundant label text for VoiceOver is removed from the exit node picker.

iOS

- [Taildrop](#) supported nodes are shown in Device Details.
- Redundant label text for VoiceOver is removed from the exit node picker.

Nov 25, 2025

Tailscale v1.90.9

[Update instructions](#) →

All platforms

- [tailscaled](#) no longer deadlocks during event bursts.
- The client no longer hangs after wake up when [port mapping](#) is in use and interfaces are slow to become available.

Android

- [DNS](#) continues working when switching from cellular to Wi-Fi connections.

Nov 19, 2025

Tailscale v1.90.8

[Update instructions](#) →

Note: v1.90.7 was an internal-only release.

All platforms

- Panic issue related to [Peer Relays](#) is resolved.
- Deadlock issue no longer occurs when handling Peer Relays endpoint allocation requests.
- Memory leak in Peer Relays is resolved.

Linux

- Nodes without the [tailscaled --statedir](#) flag or the [TS_STATE_DIR](#) environment variable no longer fail to enforce signing checks in tailnets with [Tailnet Lock](#) enabled. This fix addresses a security vulnerability described in [TS-2025-008](#).

macOS

- Connectivity issue related to sleep and wake is resolved.

Oct 31, 2025

Tailscale v1.90.6

[Update instructions](#) →

App connectors

- Routes no longer stall and fail to apply when updated repeatedly in a short period of time.

Oct 30, 2025

Tailscale v1.90.5

[Update instructions](#) →

Linux

- [Tailscale SSH](#) no longer hangs for 10s when connecting to tsrecorder. This affected tailnets that use [Tailscale SSH recording](#).

Oct 28, 2025

Tailscale v1.90.4

[Update instructions](#) →

All platforms

- A deadlock issue no longer occurs in the client when checking for the network to be available.

Linux

- `tailscaled` no longer sporadically panics when a [Trusted Platform Module](#) (TPM) device is present.

Windows

- `tailscaled` no longer sporadically panics when a [Trusted Platform Module](#) (TPM) device is present.

WASM

- The JS/WASM client used by `tsconnect` no longer crashes unexpectedly.

Oct 27, 2025

Tailscale v1.90.3

[Update instructions](#) →

All platforms

- `tailscaled` shuts down as expected and without panic.

Linux

- `tailscaled` starts up as expected in a no router configuration environment.

macOS

- The Tailscale dock icon closes as expected when the client is not using the [windowed UI \(beta\)](#).

FreeBSD

- `tailscaled` starts up as expected in a no router configuration environment.

OpenBSD

- `tailscaled` starts up as expected in a no router configuration environment.

Oct 24, 2025

Tailscale v1.90.2

[Update instructions](#) →

Linux

- An [iptables](#) regression on non-amd64/arm64 platforms is resolved, and the client starts as expected.
- Running Tailscale on devices equipped with Trusted Platform Module (TPM) 1.x no longer causes the `tailscaled` daemon to fail.

Oct 23, 2025

Tailscale v1.90.1

[Update instructions](#) →

Note: 1.90.0 was a [release candidate](#) intended for testing only.

All platforms

- Clients can use configured DNS resolvers for all domains even when the client also uses an exit node using the [nameserver](#) settings in the **DNS** page of the admin console.
- **Node keys** will be renewed seamlessly, so clients will maintain existing connections while re-authenticating.
- Go is updated to version 1.25.3.
- Unnecessary path discovery packets over [DERP servers](#) are suppressed.

Linux

- Node key sealing is GA ([generally available](#)) and enabled by default. Existing nodes will migrate to node key sealing automatically on upgrade. For more information, including how to opt out, refer to [Secure node state storage](#).

Windows

- Node key sealing is GA ([generally available](#)) and enabled by default. For more information, refer to [Secure node state storage](#).

macOS

- The **Hide Dock Icon** checkbox located in **Settings** lets you remove the Tailscale icon from the macOS dock when the client window is closed.
- The [tailscale drive](#) CLI command for sharing [Tailandrive](#) directories is no longer available. Use the client GUI for sharing directories instead.
- Node key sealing is GA ([generally available](#)) and enabled by default. For more information, refer to [Secure node state storage](#).
- [Exit node](#) selection using the [macOS Shortcuts](#) app work as expected.
- Accounts displayed using the macOS menu bar Tailscale icon load as expected.
- Client users preference for automatic/recommended exit node selection is remembered as expected.

iOS

- [Exit node](#) selection using the [iOS Shortcuts](#) app work as expected.
- Client users preference for automatic/recommended exit node selection is remembered as expected.

Android

- Client is able to establish [direct connections](#) as expected.

Oct 14, 2025

Tailscale v1.88.4

[Update instructions](#) →

macOS

- The macOS Firewall system setting **Block all incoming connections** no longer causes intermittent connectivity disruptions when enabled.

Sep 25, 2025

Tailscale v1.88.3

[Update instructions](#) →

All platforms

- [Control plane](#) connection issues which might have resulted in timing out during retries.

macOS

- [Taildrive](#) list of devices loads as expected when selecting **File Sharing > Choose Shared Folders**.

iOS

- The UI and device list display as expected when initially connecting to the tailnet.

OpenBSD

- The client starts as expected when using the `tailscale up` command for the first time or re-authenticating a node.

Sep 17, 2025

Tailscale v1.88.2

[Update instructions](#) →

macOS

- The **Settings** button displays correctly when no account is logged in to the client.
- [UserDefaults](#), which apps and [system policies](#) use to store and read preferences, force string values like `true` or `1` into Booleans as expected.

iOS

- [UserDefaults](#), which apps and [system policies](#) use to store and read preferences, force string values like `true` or `1` into Booleans as expected.

Sep 11, 2025

Tailscale v1.88.1

[Update instructions](#) →

Note: v1.88.0 was an internal-only release.

All platforms

- [Tailscale CLI](#) prompts users to confirm with *y/n* before proceeding with impactful actions.
- Go is updated to version 1.25.1.
- [Tailscale SSH](#) works as expected when using an IP address instead of a hostname and [MagicDNS](#) is disabled.
- [Taildrive](#) folder sharing works correctly even when the `su` command is not present on the Linux or other Unix-like host.
- Taildrive files remain consistently accessible.

Linux

- The [system tray application](#) for Linux desktops can be enabled to display some of the GUI options available in other Tailscale clients, including [fast user switching](#) and [exit node](#) selection.

Windows

- The existing [ExitNodeID=auto:any](#) system policy supports the new `ExitNode.AllowOverride` policy option that lets users select a different exit node while still requiring exit node usage.

macOS

- The existing [ExitNodeID=auto:any](#) system policy supports the new `ExitNode.AllowOverride` policy option that lets users select a different exit node while still requiring exit node usage.
- Windowed UI mode ([beta](#)) provides an updated client experience. To test, go to the **Settings** page of the admin console and toggle **Redesigned macOS Client UI**. Once enabled, all macOS clients display the new interface.
- `UseSystemProxy` default setting to indicate whether Tailscale respects proxy settings defined in **System Settings**.
- [advertiseExitNode](#) system policy is available on macOS.
- macOS 12 is the minimum supported version.
- Automatic recommended exit node selection.
- UI improvements for iOS 26 and macOS 26 compatibility.

iOS

- UI improvements for iOS 26 and macOS 26 compatibility.

QNAP

- New [QNAP](#) builds are available again. At the time of this release, you can manually download the update from our [packages site](#). After a period of time, the update will also be available in [QNAP App Center](#).

Aug 7, 2025

Tailscale v1.86.4

[Update instructions](#) →

Note: v1.86.3 was an internal-only release.

macOS

- [EncryptState](#) system policy changes are applied without needing to restart the [system extension](#).
- Startup crash on a fresh install of the [Standalone variant](#) of the client when the EncryptState system policy is enabled.

Android

- Persistent notifications about the [Taildrop](#) directory picker. The notification only displays on the first attempt to use the feature.

Jul 29, 2025

Tailscale v1.86.2

[Update instructions](#) →

Note: v1.86.1 was an internal-only release.

All platforms

- A deadlock issue that may have occurred in the client.
- An occasional crash when establishing a new port mapping with a gateway or firewall.

macOS

- Issue preventing the reading of existing state files that may have required device re-approval if [device approval](#) is enabled on the tailnet.
- A spurious warning about uninstalling the system extension when upgrading the client.
- `tailscale syspolicy` CLI command output displays correctly when the `KeyExpirationNotice` or `ReconnectAfter` [system policies](#) are configured.

Windows

- `tailscale syspolicy` CLI command output displays correctly when the `KeyExpirationNotice` or `ReconnectAfter` [system policies](#) are configured.

Jul 24, 2025

Tailscale v1.86.0

[Update instructions](#) →

Note: Tailscale halted the rollout of version 1.86.0 for macOS on July 25, 2025, and for all other platforms on July 28, 2025, due to multiple regressions.

All platforms

- [tsStateEncrypted](#) device posture attribute for checking whether the Tailscale client state is encrypted at rest.
- [Cross-site request forgery \(CSRF\)](#) issue that may have resulted in a log in error when accessing the [web interface](#).
- Hostnames are verified as expected when using CONNECT HTTPS proxy to connect to the [control plane](#).
- [Recommended exit node](#) when the previously recommended exit node is offline.

Linux

- `tailscale up --exit-node=auto:any` and `tailscale set --exit-node=auto:any` CLI commands track the [recommended exit node](#) and automatically switches to it when available exit nodes or network conditions change.
- `tailscaled` CLI command flag `--encrypt-state` encrypts the node state file on the disk using [trusted platform module \(TPM\)](#).

Windows

- `tailscale up --exit-node=auto:any` and `tailscale set --exit-node=auto:any` CLI commands track the [recommended exit node](#) and automatically switches to it when available exit nodes or network conditions change.
- [EncryptState](#) system policy enforces storing the node state file in encrypted format on disk using [trusted platform module \(TPM\)](#).
- Selecting **Recommended** from the exit node picker makes the Tailscale client track the [recommended exit node](#) and automatically switch to it when available exit nodes or network conditions change.
- [AlwaysOn](#) system policy is enforced as expected.
- System tray icon display a notification when the selected exit node is unavailable.
- [Mullvad exit node](#) picker hides after switching from a profile with Mullvad exit nodes to one without any exit nodes.
- WDAP/PAC proxy detection on Windows 10 1607 and earlier to ensure successful connectivity when a proxy is required.

macOS

- `tailscale up --exit-node=auto:any` and `tailscale set --exit-node=auto:any` CLI commands track the [recommended exit node](#) and automatically switches to it when available exit nodes or network conditions change.
- [ReconnectAfter](#) system policy setting, which configures the maximum period of time between a user disconnecting Tailscale and the client automatically reconnecting.
- [EncryptState](#) system policy enforces storing the node state file in the [Keychain](#). The App Store variant of the client always uses the Keychain regardless of this setting.
- [OnboardingFlow](#) system policy enforces the suppression of the onboarding flow that displays when the client is installed. This replaces the deprecated [TailscaleOnboardingSeen](#) system policy.
- **Remove all accounts** option in the [Debug](#) menu.
- [TailscaleOnboardingSeen](#) system policy is deprecated. Use the new [OnboardingFlow](#) system policy instead.
- Selecting **Recommended** from the exit node picker makes the Tailscale client track the [recommended exit node](#) and automatically switch to it when available exit nodes or network conditions change.
- [AlwaysOn](#) system policy is enforced as expected.
- [Shortcut](#) action issues.

iOS

- Selecting **Recommended** from the exit node picker makes the Tailscale client track the [recommended exit node](#) and automatically switches to it when available exit nodes or network conditions change.
- **Reset keychain** option issues.
- [Shortcut](#) action issues.
- [Taildrop](#) resending issues.

tvOS

- Selecting **Recommended** from the exit node picker makes the Tailscale client track the [recommended exit node](#) and automatically switch to it when available exit nodes or network conditions change.

Jun 26, 2025

Tailscale v1.84.3

[Update instructions](#) →

Note: The Tailscale v1.84.3 client release includes fixes for Android TV only, and is exclusively released for Android TV.

Android TV

- Internal issue.

Jun 9, 2025

Tailscale v1.84.2

[Update instructions](#) →

Windows

- This release is signed with a new code signing certificate. The certificate subject and issuer remain unchanged, but the certificate has a new serial number.

May 29, 2025

Tailscale v1.84.1

[Update instructions](#) →

macOS

- DNS drops when changing networks.

iOS

- Setting to toggle subnet routing.

- Issue where Taildrop notifications may not be presented.
- Issue where subnet routing would default to off.

Android

- Issue where Mullvad nodes may be listed as tailnet devices.
- Issue where subnet routing would default to off.
- Present a modal dialog that explains why you are prompted to select a directory.

May 21, 2025

Tailscale v1.84.0

[Update instructions](#) →

All platforms

- The `--reason` flag is added to the `tailscale down` command.
- [Tailscale CLI](#) commands throw an error if multiple of the same flag are detected.
- Network connectivity issues when creating a new profile or switching profiles while using an exit node.

Linux

- DNS-over-TCP fallback works correctly with upstream servers reachable only via the tailnet.

Windows

- [AlwaysOn.Enabled](#) and [AlwaysOn.OverrideWithReason](#) policy settings, which enable and configure a Tailscale client mode where the client stays connected at all times, unless an exception applies.
- [ReconnectAfter](#) policy setting, which configures the maximum period of time between a user disconnecting Tailscale and the client automatically reconnecting.
- When **Always On** mode is enabled, Tailscale connects as soon as a user signs in to the device and stays connected, regardless of whether the GUI is running. This enables access to tailnet resources, such as network-mapped drives, earlier in the sign-in process, and can also be used on headless Windows environments.
- [EnableDNSRegistration](#) policy setting, which configures whether Tailscale IP addresses should be registered with Active Directory DNS.
- The Tailscale GUI starts for all signed-in users when the client is installed.
- DNS-over-TCP fallback works correctly with upstream servers reachable only via the tailnet.
- Issue where the Tailscale GUI would not start if the client was installed via Group Policy or mobile device management (MDM) while a user was already signed in.
- Issue where the Tailscale GUI did not auto-start after a client update.

macOS

- [AlwaysOn.Enabled](#) and [AlwaysOn.OverrideWithReason](#) policy settings, which enable and configure a Tailscale client mode where the client stays connected at all times, unless an exception applies.
- `ForceEnabled` policy setting is deprecated in favor of the `AlwaysOn` policy setting.
- DNS-over-TCP fallback works correctly with upstream servers reachable only via the tailnet.
- Tailscale automatically recreates and/or reactivates its VPN configuration on start.

- Occasional crash in client during engine updates.
- [Taildrop](#) share sheet displays the correct error page when the tunnel is not connected.
- Hostname detection is improved for macOS clients running on macOS v15.x.
- Client (GUI) logs are properly captured and recorded in bug reports.

iOS

- [AlwaysOn.Enabled](#) and [AlwaysOn.OverrideWithReason](#) policy settings, which enable and configure a Tailscale client mode where the client stays connected at all times, unless an exception applies.
- [ForceEnabled](#) policy setting is deprecated in favor of the [AlwaysOn](#) policy setting.
- [Taildrop](#) share sheet displays the correct error page when the tunnel is not connected.
- Tailscale automatically recreates and/or reactivates its VPN configuration on start.
- Client (GUI) logs are properly captured and recorded in bug reports.
- Occasional crash in client during engine updates.

tvOS

- Tailscale automatically recreates and/or reactivates its VPN configuration on start.
- Client (GUI) logs are properly captured and recorded in bug reports.
- Occasional crash in client during engine updates.

Android

- [ReconnectAfter](#) policy setting, which configures the maximum period of time between a user disconnecting Tailscale and the client automatically reconnecting.
- Issue where Tailscale was disconnecting after excluding apps via split tunneling.

Apr 17, 2025

Tailscale v1.82.5

[Update instructions](#) →

All platforms

- A panic issue related to [CUBIC congestion control](#) in [userspace mode](#) is resolved.

macOS

- The VPN approval message during the client installation displays as expected.
- An issue related to the reachability of upstream DNS servers with loopback IPs is resolved.

Windows

- A service panic issue on the 32-bit version of Windows 10 is resolved.

Apr 15, 2025

Tailscale v1.82.4

[Update instructions](#) →

Note: Tailscale v1.82.4 includes fixes for Android devices only, and is exclusively released for Android. Tailscale v1.82.2 and v1.82.3 were internal-only releases.

Android

- An issue that might have resulted in the Tailscale app crashing on devices running versions earlier than Android 13 is resolved.

Mar 27, 2025

Tailscale v1.82.1

[Update instructions](#) →

Note: v1.82.1 includes fixes for Android devices only, and is exclusively released for Android.

Android

- Device search is available on Android TV running Android 13 or later.
- Enhanced device search UI is available on all devices running Android 13 or later.

Mar 26, 2025

Tailscale v1.82.0

[Update instructions](#) →

All platforms

- [DERP](#) functionality within the client supports certificate pinning for self-signed IP address certificates for those unable to use [Let's Encrypt](#) or WebPKI certificates.
- Go is updated to version 1.24.1
- NAT traversal code uses the DERP connection that a packet arrived on as an ultimate fallback route if no other information is available, in the event of a slow or misbehaving server.
- [Captive portal](#) detection reliability is improved on some in-flight Wi-Fi networks, including British Airways and WestJet.
- Port mapping success rate is improved by retrying in additional error cases.
- [Web interface](#) setting changes occur as expected and without error.

macOS

- The [.pkg installer](#) size is decreased by 35%.
- Memory leak issue related to [shortcuts](#) is resolved.
- [MagicDNS](#) intermittent configuration failures no longer occur when waking from sleep.
- Seamless [key renewals](#) occur as expected, ensuring the client remains connected.

iOS

- Memory leak issue related to [shortcuts](#) is resolved.
- [MagicDNS](#) intermittent configuration failures no longer occur when waking from sleep.

Android

Note: The Android client release for v1.82.0 was delayed and moved into the v1.82.1 client release instead.

App connectors

- [Port mapping](#) success rates for [app connectors](#) are improved.

Mar 4, 2025

Tailscale v1.80.3

[Update instructions](#) →

Linux

- [Web interface](#) setting changes occur as expected and without error.

App connectors

- [App connectors](#) respond to DNS queries and update routes without failure. Previously, DNS resolution failures may have occurred due to a routing deadlock issue.

Feb 13, 2025

Tailscale v1.80.2

[Update instructions](#) →

All platforms

- Nodes could lose the display names of owners of peers in rare cases. This had manifested in missing names in [tailscale status](#) and could prevent incoming [Tailscale SSH](#) connections from being accepted. The behavior is reverted to that of v1.78.x and earlier.

Linux

- SSH clients that skip the [none](#) auth method and immediately try [publickey](#) can connect to [Tailscale SSH](#) as expected. The behavior is reverted to that of v1.78.x and earlier.

macOS

- SSH clients that skip the none auth method and immediately try publickey can connect to Tailscale SSH as expected. The behavior is reverted to that of v1.78.x and earlier.

FreeBSD

- SSH clients that skip the none auth method and immediately try publickey can connect to Tailscale SSH as expected. The behavior is reverted to that of v1.78.x and earlier.

Feb 6, 2025

Tailscale v1.80.1

[Update instructions](#) →

macOS

- System extension uninstalled message no longer appears erroneously when removing third-party system extensions while Tailscale is running.
- Resolved an issue that could have caused the network extension to crash in rare cases while parsing the macOS routing table.

iOS

- Resolved an issue that could have caused the network extension to crash in rare cases while parsing the iOS routing table.

tvOS

- Resolved an issue that could have caused the network extension to crash in rare cases while parsing the tvOS routing table.

Jan 30, 2025

Tailscale v1.80.0

[Update instructions](#) →

All platforms

- [Hostname](#) system policy is added for overriding the device hostname configured by the operating system, using an MDM solution.
- [tailscale configure](#) CLI command and corresponding subcommands are no longer in [alpha](#), except for the subcommand `kubeconfig`, which remains in alpha.
- [Web interface](#) displays a **Login** button instead of the **Reauthenticate** button when adding a new device to your tailnet.
- [Tailscale Funnel](#) configuration on devices displays errors when [incoming connections](#) are not permitted and connections are disallowed.
- Connections to a [custom coordination server](#) that does not support HTTPS will no longer fail when a custom port number is specified.

Linux

- TLS certificate requests from [Let's Encrypt](#) include the device's DNS name in the CSR's SAN extension and set the Common Name field.
- [Tailscale Funnel](#) disabled on a device no longer displays enabled in the admin console.

Windows

- Onboarding flow is added for easier initial setup of the app.
- TLS certificate requests from [Let's Encrypt](#) include the device's DNS name in the CSR's SAN extension and set the Common Name field.
- Client installs as expected when using [Group Policy Software Installation](#) (GPSI).
- Race conditions that result in an incorrect state or a deadlock no longer cause issues when multiple Windows users are logged in simultaneously.

macOS

- [configure sysex](#), [configure sysex deactivate](#), and [configure sysex status](#) CLI commands are added to the Standalone variant for managing the activation flow of the macOS system extension programmatically.
- Standalone variant detects if the system extension is manually disabled or uninstalled by the user and displays a notice in the client UI.
- **Flush DNS Cache** option is added to the [Debug menu](#).
- TLS certificate requests from [Let's Encrypt](#) include the device's DNS name in the CSR's SAN extension and set the Common Name field.
- App preferences re-set configures **Use Tailscale Subnets** to On and **Allow Incoming Connections** to Off as these are the default settings.
- [Find Devices](#) shortcut action no longer hangs.
- Standalone variant works as expected when users are not members of `staff` macOS user group.

iOS

- [Auth keys](#) can be used for connecting to a [custom coordination server](#).
- VPN extension no longer runs when logging out.
- [Find Devices](#) shortcut action no longer hangs.

tvOS

- [Auth keys](#) are supported for [authenticating an Apple TV](#) in your tailnet.
- Auth keys can be used for connecting to a [custom coordination server](#).
- VPN extension no longer runs when logging out.

Android

- Devices can be configured as a [subnet router](#) in the **Settings** menu of the app.

Dec 13, 2024

Tailscale v1.78.3

[Update instructions →](#)

Note: Tailscale v1.78.2 was an internal-only release.

Containers

- Unit test that would previously fail if run in a container.

iOS

- **Advanced DNS Settings** view unexpectedly dismissed on iPhone.

Android

- Work in progress search bar is hidden behind a flag until the feature is ready.

Dec 5, 2024

Tailscale v1.78.1

[Update instructions →](#)

All platforms

- Issue which resulted in an unwanted change in source code line endings.

Tailscale v1.78.0

[Update instructions →](#)

All platforms

- [Client metrics](#) have been added, to provide insights into Tailscale client behavior, health, and performance.
- [tailscale metrics](#) command has been added, to expose and collect client metrics for use with third-party monitoring systems.
- [tailscale syspolicy](#) command has been added, to list [system policies](#), reload system policies, or view errors related to the system policies configured on the device.
- Tailscale system policies are applied immediately when pushed via mobile device management (MDM) or Group Policy, without requiring a client restart.
- [Tailscale SSH session recording](#) detects the disappearance of the recorder node sooner. This fix addresses a security vulnerability described in [TS-2024-013](#).

Windows

- UI customization [system policies](#) are configurable for both devices and users.

macOS

- UI to configure custom DNS servers to use for Tailscale-bound traffic when Tailscale DNS is disabled in settings.
- The macOS configuration report diagnostic tool can collect a larger amount of diagnostics when requested by Tailscale support. This includes system and process logs on the [Standalone variant](#).
- **Update Available** notifications include a link to the client changelog.
- On macOS Sequoia, in **System Settings.app > Login Items & Extension**, Tailscale is listed as **Tailscale Network Extension** instead of **IPNExtension**, to reduce user confusion.
- Performance optimizations reduce CPU and memory usage when parsing network maps, especially for users on larger and busy tailnets.
- Performance optimizations at the UI layer reduce flickering of the menus, especially for users on larger and busy tailnets where the contents of the network map change very frequently.
- Error messages displayed when failing to toggle a setting are improved and easier to understand.

iOS

- UI to configure custom DNS servers to use for Tailscale-bound traffic when Tailscale DNS is disabled in settings.
- On iPhones and iPads running iOS 18, the VPN can be toggled from Control Center. Hold down in an empty space to add the Tailscale Control.

tvOS

- UI to configure custom DNS servers to use for Tailscale-bound traffic when Tailscale DNS is disabled in settings.

Android

- Authentication by using a [generated code](#) is available for Android TV users.
- Search bar shows suggestions.
- The default avatar displays if the user has no profile picture.
- False positive health warnings in the UI are reduced.
- Health warnings are no longer displayed in the UI after stopping Tailscale.
- Crashes when sharing a file using [Taildrop](#) from another Android app are reduced.
- UI padding of the main app toolbar is improved.

Nov 6, 2024

Tailscale v1.76.6

[Update instructions](#) →

Note: v1.76.4 and v1.76.5 were internal-only releases.

All platforms

- Logging for when clients move home [DERP](#) regions is improved.
- Tailscale clients no longer move their home DERP server prematurely in response to unusual latency at very specific times.

Android

- Android app no longer terminates unexpectedly when performing network transitions.

Oct 21, 2024

Tailscale v1.76.3

[Update instructions](#) →

Note: v1.76.3 includes fixes for Windows devices only, and is exclusively released for Windows.

Windows

- Mullvad VPN submenu no longer fails to populate with [Mullvad exit nodes](#) if there aren't any non-Mullvad exit nodes in the tailnet.

Oct 17, 2024

Tailscale v1.76.2

[Update instructions](#) →

Note: v1.76.2 includes fixes for Android TV devices only, and is exclusively released for Android.

Android

- D-Pad navigation is optimized in the Tailscale app on Android TV devices.

Oct 16, 2024

Tailscale v1.76.1

[Update instructions](#) →

All platforms

- [tailscale netcheck](#) CLI command no longer crashes when performing diagnostics on networks lacking UDP connectivity.
- Improperly formatted SERVFAIL responses no longer cause DNS timeouts when using an [exit node](#).

Linux

- dbus login sessions no longer fail on systems where `/bin/login` is missing.

Android

- Android application no longer crashes in certain configurations when editing the [app-based split tunneling](#) settings.

Oct 10, 2024

Tailscale v1.76.0

[Update instructions](#) →

All platforms

- Clients lacking UDP connectivity no longer skip performing fallback latency measurements with [DERP servers](#).
- Warnings no longer display unnecessarily.
- Tailscale connectivity on flights using inflight internet Wi-Fi (such as Alaska Airlines) no longer fails.
- Service-related processes no longer run unnecessarily when services are disabled on the tailnet.
- Error messages include explanations in addition to the HTTP status code.

Linux

- [Tailscale SSH](#) supports sending environment variables to hosts. It's also possible to specify permitted environment variables using the `acceptEnv` field.
- [Tailscale SSH](#) no longer breaks some terminal applications by omitting pixel width and height when resizing the application window.

Windows

- Ping messages sent through [subnet routers](#) to unreachable hosts no longer generate ping responses.

macOS

- [Tailscale SSH](#) supports sending environment variables to hosts. You must specify permitted environment variables using the `acceptEnv` field.
- Tailscale `.pkg` installer for the [standalone variant](#) prevents potential conflicts by showing a warning if it detects a Homebrew install of Tailscale.
- [Bug report](#) view shows a warning if Tailscale detects that Cloudflare WARP is installed. Some Cloudflare WARP configurations conflict with Tailscale.
- DNS settings no longer improperly set when keys expire or Tailscale stops.

iOS

- Battery usage is improved when [MagicDNS](#) is enabled. The improvement comes from adjusting the timeout of [DNS over HTTPS \(DoH\)](#) for idle connections and requiring a TLS 1.3 handshake when establishing a connection with the DoH server.
- DNS settings no longer improperly set when keys expire or Tailscale stops.

tvOS

- DNS settings no longer improperly set when keys expire or Tailscale stops.

Android

- [Account switcher](#) displays the server hostname if the account uses a custom coordination server.
- Battery usage is improved when [MagicDNS](#) is enabled. The improvement comes from adjusting the timeout of [DNS over HTTPS \(DoH\)](#) for idle connections and requiring a TLS 1.3 handshake when establishing a connection with the DoH server.
- Quick tile toggle no longer fails to turn on Tailscale if Tailscale had been manually disconnected before it was last shut down.

Oct 2, 2024

Tailscale v1.74.2

[Update instructions](#) →

Tailscale v1.74.2 addresses an issue for iOS, and is exclusively released for that platform.

iOS

- The Tailscale app launches as expected when **Wi-Fi Calling on This iPhone** is enabled in the iOS **Cellular** settings.

Sep 18, 2024

Tailscale v1.74.1

[Update instructions](#) →

Tailscale v1.74.1 addresses issues for Linux and Android, and is exclusively released for those platforms.

Linux

- Linux-only NAT traversal optimization added in v1.74.0 is now disabled following a bug report. The behavior is reverted to that of v1.72.x and earlier and will be re-added in a future release.

Android

Note: The Android client release for v1.74.0 was delayed and moved into the v1.74.1 client release instead.

- Device network change detection is improved to reflect accurate [Tailscale DNS](#) configuration updates.
- [System policies](#) for the Android client on ChromeOS work as expected.

Sep 12, 2024

Tailscale v1.74.0

[Update instructions](#) →

All platforms

- [AuthKey](#) system policy can be used to authenticate a device with Tailscale using an [MDM solution](#).
- `tailscale dns` CLI command is added for accessing [Tailscale DNS](#) settings and status.
- Go is updated to version 1.23.1.
- [Tailnet Lock](#) long rotation signatures are truncated automatically to avoid excessive growth.
- **Log In** option in the client works as expected.

Linux

- [TCP generic receive offload](#) (GRO) support is added for improved userspace mode throughput.
- [TCP generic segmentation offload](#) (GSO) is re-introduced for supporting improved userspace mode throughput. This was initially introduced in Tailscale v1.72.0 and then rolled back in v1.72.1.

Windows

- The client no longer connects to a tailnet automatically when restarting or switching profiles.
- Profiles created as Local System with Unattended Mode enabled are retained after a reboot.

macOS

- The [open-source variant](#) of the Tailscale client can now read the system [DNS configuration](#) to provide DNS resolution when `tailscale set --accept-dns` or `tailscale up --accept-dns` is enabled and the **Override local DNS** option in the **DNS** page of the admin console is disabled.
- DNS resolution continues to work after a key expires.

tvOS

- The [ping](#) feature allows you to observe connectivity performance between your Apple TV and other devices in your tailnet.

Android

Note: The Android client release for v1.74.0 was delayed and moved into the v1.74.1 client release instead.

- [Tailscale DNS](#) works as expected when switching between Wi-Fi and cellular networks.
- [System policies](#) for the Android client on ChromeOS work as expected.

Aug 26, 2024

Tailscale v1.72.2

[Update instructions](#) →

Tailscale v1.72.2 addresses issues for macOS, iOS, and tvOS, and is exclusively released for those platforms.

macOS

- An issue that could trigger a VPN permission prompt when starting Tailscale while another VPN app was already active is fixed.
- An issue that could prevent Tailscale from automatically launching at login on some Macs is fixed.

iOS

- An issue that could trigger a VPN permission prompt when starting Tailscale while another VPN app was already active is fixed.

tvOS

- An issue that could trigger a VPN permission prompt when starting Tailscale while another VPN app was already active is fixed.

Aug 22, 2024

Tailscale v1.72.1

[Update instructions](#) →

Tailscale v1.72.1 addresses a Linux-specific issue, and is exclusively released for the Linux platform and containers.

Linux

- TCP generic segmentation offload (GSO) support for userspace mode is removed.
- DNS over TCP failures when querying the Tailscale-internal resolver are fixed.

Aug 19, 2024

Tailscale v1.72.0

[Update instructions](#) →

All platforms

- [Captive portal detection](#) is now supported.
- The `tailscale cert` command now contains the `--min-validity` flag. Use this flag to request a specified minimum remaining validity on the returned certificate. This flag is intended for automation, like cron jobs, that periodically refreshes certificates.
- The `tailscale lock` command now supports passing keys as files. To pass a key as a file, use the prefix `file:` followed by the path to the file: `file:<path-to-key-file>`.
- A health warning is now raised if Tailscale is unable to forward DNS queries to the configured resolvers.
- An increase in send and receive buffer sizes for userspace mode TCP improves throughput over high latency paths.

Linux

- The addition of TCP generic segmentation offload (GSO) support to userspace mode improves throughput.

macOS

Note: macOS 10.15 Catalina is no longer supported. See the [v1.60.0 changelog](#) for our initial end of life announcement.

- Notifications are sent when a captive portal is detected.
- Health warnings in the UI are now sorted by their severity level.
- Reliability of the authentication process when launching the web browser is improved.
- The VPN tunnel is no longer automatically restarted if toggling Tailscale from the system VPN settings without disabling [VPN On Demand](#) first.

iOS

- Notifications are sent when a captive portal is detected.
- Health warnings are displayed when connectivity is impacted.
- An error message is displayed while attempting to start the VPN when both Wi-Fi and cellular interfaces are down, instead of failing silently.
- The VPN tunnel is no longer automatically restarted if toggling Tailscale from the system VPN settings without disabling [VPN On Demand](#) first.

tvOS

- Notifications are sent when a captive portal is detected.
- The VPN tunnel is no longer automatically restarted if toggling Tailscale from the system VPN settings without disabling [VPN On Demand](#) first.

Android

- Health warnings, if any are present, are displayed in the main view of the app.

Jul 17, 2024

Tailscale v1.70.0

[Update instructions](#) →

All platforms

- Restrict [recommended](#) and automatically selected exit nodes using the new `AllowedSuggestedExitNodes` [system policy](#). Applies only to platforms that support [system policies](#).
- Improved [NAT traversal](#) for some uncommon scenarios.
- Optimized [sending firewall rules to clients](#) more efficiently.
- [Exit node suggestion](#) CLI command now prints the hostname (which you can use with the `tailscale set` command).
- [Taildrive](#) share paths configured through the CLI resolve relative to where you run the `tailscale` command.

Linux

- Switching from unstable to stable tracks using the `tailscale update` command now works correctly.

Windows

- Use the value `auto:any` to automatically select an [exit node](#) for the existing `ExitNodeID` [system policy](#). Available for [Enterprise plan](#) users only.
- The new `AllowedSuggestedExitNodes` [system policy](#) restricts which exit nodes Tailscale [recommends](#) or automatically selects.
- DNS leak issue.
- Switching from unstable to stable tracks using the `tailscale update` command now works correctly.
- [Taildrive](#) server no longer starts unnecessarily when no drives are configured.

macOS

Note: As previously announced, Tailscale v1.70 is the last version to support macOS 10.15 Catalina. macOS 10.15 is no longer supported by Apple and no longer receives security updates. Users still running macOS 10.15 should update to a newer version of macOS to continue receiving security updates and new features.

- Toggle Tailscale DNS from Siri or the Shortcuts app.
- Receive health notifications in the client menu on macOS to inform you about lack of internet connectivity, firewalls blocking Tailscale, misconfiguration issues, and other issues. Health issues that affect [connectivity](#) also change the Tailscale icon in the system menubar to show an exclamation mark.
- On MacBooks with a notch in the display, a notification window will now appear if the Tailscale icon is hidden behind the notch due to too many menubar items.
- The Tailscale client now warns you when the macOS built-in [content filter \(Screen Time\)](#) prevents Tailscale from connecting.
- Use the value `auto:any` to automatically select an exit node for the existing `ExitNodeID` [system policy](#). Available for [Enterprise plan](#) users only.
- The exit node picker no longer presents exit node suggestions if the organization enforces always using the suggested exit node using the `ExitNodeID` [system policy](#).
- Disconnect shortcut no longer connects to the VPN tunnel if executed when Tailscale is disconnected.
- [Taildrive](#) server no longer starts unnecessarily when no drives are configured.
- Increased the reliability of the **Install Updates Automatically** setting.

iOS

- Toggle Tailscale DNS from Siri or the Shortcuts app.
- Use the value `auto:any` to automatically select an exit node for the existing `ExitNodeID` [system policy](#). Available for [Enterprise plan](#) users only.
- [wireguard-go](#) memory pool deadlock issue is resolved.
- Disconnect shortcut no longer connects to the VPN tunnel if executed when Tailscale is disconnected.
- User interface no longer flickers when selecting an exit node.

tvOS

- Use the value `auto:any` to automatically select an exit node for the existing `ExitNodeID` [system policy](#). Available for [Enterprise plan](#) users only.
- [wireguard-go](#) memory pool deadlock issue is resolved.
- User interface no longer flickers when selecting an exit node.

Android

- Access ping information and connection status by long-pressing on a device in the devices list and selecting **Ping**.
- Use [split tunneling](#) to force or exclude app traffic through your tailnet.
- [wireguard-go](#) memory pool deadlock issue is resolved.

Jul 2, 2024

Tailscale v1.68.2

[Update instructions](#) →

All Platforms

- [Tailnet Lock](#) validation of rotation signatures now permits multiple nodes signed by the same pre-signed reusable auth key.

macOS

- Wake from sleep reliability is improved for re-connections and transitions between networks.

iOS

- Wake from sleep reliability is improved for re-connections and transitions between networks.

Jun 14, 2024

Tailscale v1.68.1

[Update instructions](#) →

All Platforms

- [4via6 subnet router](#) advertisement works as expected.

Linux

- [Tailscale SSH](#) access to Security-Enhanced Linux (SELinux) machines works as expected.

Android

- Android TV navigation is improved.

Jun 12, 2024

Tailscale v1.68.0

[Update instructions](#) →

All Platforms

- [Auto-updates](#) are available for containers. The tailnet-wide default is ignored in containers.

- When enabled, auto-updates get applied even if the node is down or disconnected from the [coordination server](#).
- [tailscale lock status](#) now prints the node's signature.
- Go is updated to version 1.22.4.

Windows

- [.exe installer](#) no longer downloads MSI packages for Windows 7 and Windows 8, automatically. See the [v1.42.0 changelog](#) for our initial end of life announcement.

macOS

- [Standalone](#) variant of the client can now install a launcher for the [Tailscale CLI](#) in `/usr/local/bin` by going to **Settings**, **CLI integration**, then **Show me how**.
- Standalone variant of the client now supports notifications when a file is received using [Tailable](#).
- Pop-up notification displays when a network might be vulnerable to a potential TunnelVision attack. For more information, see [TunnelVision vulnerability and Tailscale](#).
- Client starts up more reliably if another VPN app is running when Tailscale is enabled.
- [.pkg installer](#) terminates pre-existing copies of Tailscale and the VPN extension before proceeding with installation if Tailscale was already installed.
- TunnelBear installation is properly detected, and warns the user about incompatibility.
- Using `Exit Node` label no longer appears incorrectly in the app menu before completing onboarding, upon the first time app launch.
- Fixed a bug with split DNS domains being used as search domains after a network change.

iOS

- Battery life is optimized by offloading DNS resolution to iOS in more cases.
- Client now starts more reliably if another VPN app is running when Tailscale is enabled.
- Bug report view no longer copies the bug report ID to the clipboard automatically.
- **Reauthenticate** button for in-app key expiry notifications works as expected.
- Dark mode contains minor changes to UI colors.
- Fixed a bug with split DNS domains being used as search domains after a network change.

tvOS

- Client now starts more reliably if another VPN app is running when Tailscale is enabled.
- **Reauthenticate** button for in-app key expiry notifications works as expected.

Android

- On-off toggle state better matches the actual client state.
- Status notifications when Tailscale is disconnected are now background notifications, and tapping on notifications launches the Tailscale app.
- Client starts automatically after the first login.
- [System policy](#) (MDM) support is added for mandatory exit nodes.
- Organization name is now rendered properly when set in the [ManagedByOrganizationName](#) system policy.
- Crashing no longer occurs when launching Tailscale and another VPN application was already running.
- [Running an exit node](#) no longer lets you [use another device as an exit node](#) and vice versa.
- Home screen shows the selected exit node country and city when using [Mullvad exit nodes](#).

Note: The Tailscale client releases for containers such as the [Kubernetes operator](#), [Docker](#) image, and [tsrecorder](#) are typically released a few days after the initial client release. A separate changelog will be published when client updates for containers are available.

May 20, 2024

Tailscale v1.66.4

[Update instructions](#) →

All platforms

- Restored UDP connectivity through [Mullvad exit nodes](#).

Linux

- Stateful filtering is now off by default. Stateful filtering was introduced in 1.66.0 as a mitigation for a vulnerability described in [TS-2024-005](#), and inadvertently broke DNS resolution from containers running on the host. Most vulnerable setups are protected by other mitigations already, except when `autogroup:danger-all` is used in [ACLs](#).

May 15, 2024

Tailscale v1.66.3

[Update instructions](#) →

Note: Tailscale v1.66.2 was an internal-only release.

All platforms

- Login URLs did not always appear in the console when running `tailscale up`.

Android

- Reintroduced the Quick Settings title that v1.66.0 temporarily removed.
- Improved the VPN service connection logic, especially when rebooting the device with Always-On VPN enabled.
- The persistent VPN status notification now informs the user with a muted icon when the VPN is disconnected. VPN status notifications can be disabled in the system notification settings.
- The "Enable" button in the exit node selector banner now renders with the correct background color.

Kubernetes operator

- Starting with v1.66, the Kubernetes operator must always run the same or later version as the proxies it manages.

- [Expose cloud services](#) on cluster network to the tailnet, using Kubernetes ExternalName Services. This allows exposing cloud services, such as RDS instances, to tailnet by their DNS names.
- Expose tailnet services that use [Tailscale HTTPS](#) to cluster workloads. Refer to [#11019](#).
- Cluster workloads can now refer to Tailscale Ingress resources by their MagicDNS names. Refer to [#11019](#).
- Configure environment variables for Tailscale Kubernetes operator proxies using ProxyClass CRD. Refer to [ProxyClass API](#).
- Expose tailscaled metrics endpoint for Tailscale Kubernetes operator proxies through ProxyClass CRD. Note that the tailscaled metrics are unstable and will likely change in the future. Refer to [ProxyClass API](#).
- Configure labels for the Kubernetes operator Pods with Helm chart values. Refer to [Helm chart values](#).
- Configure affinity rules for Kubernetes operator proxy Pods with ProxyClass. Refer to [ProxyClass API](#).
- Kubernetes operator proxy init container no longer attempts to enable IPv6 forwarding on systems that don't have IPv6 module loaded. Refer to [#11867](#).

Containers

- Tailscale containers running on Kubernetes no longer error if an empty Kubernetes Secret is pre-created for the tailscaled state. Refer to [#11326](#).
- Improved the ambiguous error messages when Tailscale running on Kubernetes does not have the right permissions to perform actions against the tailscaled state Secret. Refer to [#11326](#).

May 9, 2024

Tailscale v1.66.1

[Update instructions](#) →

This release is exclusively for Linux platforms and the [standalone variant](#) of the macOS client. It is not available for other platforms.

Linux

- [tailscale set](#) command flags `--netfilter-mode`, `--snat-subnet-routes`, and `--stateful-filtering` are added.
- Issue with [nftables](#) rules for stateful filtering, introduced in v1.66.0.

macOS

- A version mismatch warning no longer displays when upgrading, if no mismatch is detected.

May 8, 2024

Tailscale v1.66.0

[Update instructions](#) →

We recommend updating all Tailscale clients to v1.66.0 or later to benefit from additional security improvements.

All platforms

- Implemented client-side quarantining for shared-in exit nodes, as a mitigation for a security vulnerability described in [TS-2024-005](#).

Linux

- Use the `--stateful-filtering` flag for the `tailscale up` to enable stateful filtering for [subnet routers](#) and [exit nodes](#), as a mitigation for a security vulnerability described in [TS-2024-005](#).

Note: This change can break existing setups that depend on forwarding connections from external hosts (internet, LAN, Docker containers, and similar) into the tailnet through a Tailscale node. If your setup depends on such forwarding, you can disable stateful filtering with the `tailscale up --stateful-filtering=false` command.

- Use [tab completion](#) to type the first few letters of a Tailscale CLI command, flag, or arguments, followed by the tab key to complete the item being typed. Set up tab completion by using the `tailscale completion` command.
- Use the `tailscale exit-node suggest` command to automatically pick an available exit node that is likely to perform best.
- [Site-to-site networking](#) now also requires `--stateful-filtering=false` in addition to `--snat-subnet-routes=false` on new subnet routers. Existing subnet routers with `--snat-subnet-routes=false` will default to `--stateful-filtering=false`.

macOS

- View a suggested [exit node](#) in the **Exit Node** picker when available.
- Generate a macOS Configuration Report `.txt` file from the **Bug Report** view to help the Tailscale support team diagnose issues.
- Improved error detection logic warns the user when a version mismatch is detected between the Tailscale client GUI and the network extension.

iOS

- See direct vs. relayed connections in the **Ping** view.
- View a suggested [exit node](#) in the **Exit Node** picker when available.
- Use [auth keys](#) to log in without using the browser.
- Search [tagged devices](#) by tag in the **Devices** list.
- Remove accounts in the **Fast User Switching** view by using a long press, without having to log out.
- Improved UI experience to log into a custom coordination server like [Headscale](#).
- The **Fast User Switching** view can now be used when Tailscale is disconnected.
- Improved error detection logic warns the user when a version mismatch is detected between the Tailscale client GUI and the network extension.
- Reduced app launch time.

tvOS

- Manage DNS configuration in the **DNS Settings** view.
- Generate a [bug report](#) identifier by navigating to **About Tailscale > Report an issue**.
- Improved error detection logic warns the user when a version mismatch is detected between the Tailscale client GUI and the network extension.

Android

We've rebuilt the Android app from the ground up, adopting a similar design that we've previously rolled out on iOS and using the latest Android best practices.

- Use new status indicators to see at-a-glance insights into node connectivity. Tap on a node to see detailed information.
- See detailed information about resolvers, domains, and routing configurations in a dedicated **DNS Settings** view.
- See the status of [Tailnet Lock](#) and node keys.
- Use [Fast user switching](#) to switch between two or more logged-in accounts on the same device, without requiring you to re-authenticate.
- Use [auth keys](#) to log in without using the browser.
- Manage Android devices in your tailnet using [Mobile Device Management](#) (MDM) solutions such as [Google Workspace](#), [Microsoft Intune](#), or [TinyMDM](#), among other tools.
- Accessibility support.
- Use dark mode as an alternative to light mode.
- The **Quick Settings** tile has been temporarily disabled, pending resolution of an issue.
- More intuitive behavior switching between exit nodes.
- Issue with LAN access during exit node use.

Apr 17, 2024

Tailscale v1.64.2

[Update instructions](#) →

Windows

- Installers are now built using WiX toolchain [version 3.14.1](#).

Synology

- DiskStation Manager UI no longer freezes for a few minutes at startup when attempting to clean unused routes. This update is applicable to the version provided on pkgs.tailscale.com.

Apr 11, 2024

Tailscale v1.64.0

[Update instructions](#) →

All platforms

- `tailscale serve` headers are now [RFC 2047](#) Q-encoded.
- Device web interface enabled by default locally on `100.100.100.100`.
- Go is updated to version 1.22.2.

macOS

- Use Tailscale for macOS as a Tailscale SSH client (Standalone variant only).
- Receive alerts when an error occurs while changing client preferences.
- Added a new [internet access policy](#) for [Little Snitch](#) users.
- The .pkg installer no longer requires a system restart after installing the client (Standalone variant only).
- Unexpected terminations for some macOS 10.15 Catalina users.
- Reduced number of alerts if the network extension terminates unexpectedly.

iOS

- Improved reliability of the ping chart presentation.

Synology

- Update certificates using the [cert](#) CLI command.
- [IPv6](#) addresses are available again.

Kubernetes operator

- [tailscale configure kubeconfig](#) now respects KUBECONFIG environment variable.
- [tailscale configure kubeconfig](#) now works with partially empty kubeconfig.
- [MSS](#) clamping for Kubernetes operator proxies using [nftables](#).

Containers

- Containers on hosts with partial support for [ip6tables](#) no longer crash.

Mar 26, 2024

Tailscale v1.62.1

[Update instructions](#) →

Linux

- Send load balancing hint HTTP request header

Windows

- Do not allow [msiexec](#) to reboot the operating system

macOS

- Issue that could cause the Tailscale system extension to not be installed upon app launch, when deploying Tailscale using MDM and using a configuration profile to pre-approve the VPN tunnel (applies to [standalone](#) variant only)

Synology

- IPv6 routing

Kubernetes operator

- [Kubernetes operator](#) proxies should not accept subnet routes

Mar 13, 2024

Tailscale v1.62.0

[Update instructions](#) →

All platforms

- [Web interface](#) now uses ACL grants to manage access on tagged devices
- [Tailscale SSH](#) connections now disable unnecessary hostname [canonicalization](#)
- [tailscale bugreport](#) command for generating diagnostic logs now contain ethtool information
- Mullvad's [family-friendly server](#) is added to the list of well known DNS over HTTPS (DoH) servers
- DNS over HTTP requests now contain a timeout
- TCP forwarding attempts in [userspace mode](#) now have a per-client limit
- Endpoints with link-local IPv6 addresses is preferred over private addresses
- WireGuard logs are less verbose
- Go is updated to version 1.22.1
- [DERP server](#) region no longer changes if connectivity to the new DERP region is degraded

Linux

- [Auto-update](#) version detection on Alpine Linux is improved
- IPv6 support detection in a container environment is improved
- DNS configuration on Amazon Linux 2023 no longer causes an infinite loop

Windows

- [ManagedByOrganizationName](#), [ManagedByCaption](#), and [ManagedByUrl](#) system policy keys are now supported
- Tailscale Tunnel WinTun adapter handling is improved
- [MSI](#) upgrades no longer ignore policy properties set during initial install

macOS

- A .pkg installer package is now available for the [standalone](#) release of the Tailscale client
- [Taildrop](#) notifications now include actions to reveal the received file in the Finder, or delete it
- [Tailnet Lock](#) settings UI displays more information about the status, including key and public key trust status
- The onboarding flow now guides the user in enabling the Tailscale system extension
- **Launch Tailscale at login** settings item can now be toggled when the Tailscale client is disconnected
- DNS behavior is improved when handling transitions between network interfaces

iOS

- Battery usage is improved
- [Taldrop](#) notifications now include actions to reveal the received file in the Files app, or delete it
- [Tailnet Lock](#) settings UI displays more information about the status, including key and public key trust status
- Unnecessary log messages are removed when triggered by changes to device power state and routing
- DNS behavior is improved when handling interface transitions between Wi-Fi and Cellular

Android

- Settings persist from previous sign-ins
- Always-on VPN handling is improved
- Custom control server is applied on first start

Kubernetes operator

- [Ingress](#) resource handling is improved when deployed before its backing Service resource
- Destination NAT (DNAT) rule management by egress proxies in [nftables](#) mode when IP address of `tailscale.com/tailnet-fqdn` changes

Feb 29, 2024

Tailscale v1.60.1

[Update instructions](#) →

All platforms

- Exposing port 8080 to other devices in your tailnet works as expected

Feb 15, 2024

Tailscale v1.60.0

[Update instructions](#) →

All platforms

- [tailscale status](#) command output now includes location-based exit nodes
- [tailscale web](#) command flag `--read-only` is added to run the web UI in read-only mode
- A warning is logged when unable to find SSH host keys
- Support added for legacy "urn:dslforum-org" port mapping services
- Build with Go 1.22
- Detect when Tailscale is running on Digital Ocean and automatically use Digital Ocean's DNS resolvers
- Expose gVisor metrics in debug mode
- Improve error message when running as non-root
- A valid login page is presented to users when attempting to log in even after leaving device unattended for several days
- An issue with noisy peer mtu discovery errors
- A potential crash when no supported port mapping services are found

Windows

- Fixed:tailscaled could be slow or cause increased CPU usage with large routing tables

macOS

Note: Tailscale v1.60.0 is built with Go 1.22 and Go 1.22 is the last release that will run on macOS 10.15 Catalina ([source](#)). We are providing notice that around August 15, 2024, Tailscale will be built with Go 1.23 at which time macOS users that want to run the latest version of Tailscale will require macOS 11 Big Sur or later. Note that macOS 10.15 Catalina is no longer supported by Apple and is no longer receiving security updates.

- New UI to add, remove, and switch between user accounts, including using custom control servers
- New UI to change client preferences
- New UI to manage updates for the Standalone variant of the client, including switching in-app between stable and unstable builds
- **VPN On Demand** is now supported on macOS, to automatically connect/disconnect Tailscale when specific conditions are triggered
- **Reset VPN Configuration** menu item in the **Debug** menu is now available to reset the system VPN configuration if needed
- An alert window is presented when the Tailscale network extension fails to start, providing suggested troubleshooting steps
- Tailscale appears in the macOS Dock when an app window is presented
- The **Network Devices** list now shows all devices known to the control server, not only those seen in the last 4 days
- The onboarding flow automatically advances once the user is connected
- A potential crash and excessive logging upon client launch
- **Start on Login** is set correctly on macOS Ventura and earlier versions

iOS

- A potential crash and excessive logging upon client launch
- Stale devices are no longer presented in the devices list

tvOS

- A potential crash and excessive logging upon client launch
- Stale devices are no longer presented in the devices list

Android

- Mullvad exit nodes now sorted to make it easier to find the best node for each location
- Mullvad tunnels are no longer shown as regular nodes in UI
- Quick settings tile now works

Synology

- An issue with stalling of SMB transfers of large files

Kubernetes operator

- A new ProxyClass custom resource that allows you to provide a custom configuration for cluster resources that the operator creates
- [ACL tags](#) for the operator can now be configured via Helm chart values
- Routing to Ingress backends that require an exact path without a slash (/) suffix

App connectors

- [App connectors](#) now flatten DNS CNAME chains down to a target A/AAAA routing record, for apps that are configured with a DNS record that is a CNAME
- Apps can be preconfigured with known routes to have those routes auto-advertised by all selected app connectors, and immediately begin to route traffic

Jan 23, 2024

Tailscale v1.58.2

[Update instructions](#) →

Note: The 1.58.1 release needed to be re-done. Use 1.58.2 instead.

All platforms

- [App connectors](#) have improved scheduling and merging of route changes under some conditions
- Crash when performing UPnP portmapping on older routers with no supported portmapping services

macOS

- Opening the **About** window no longer displays a user interface when there is no newer version

Jan 18, 2024

Tailscale v1.58.0

[Update instructions](#) →

Note: Rollout of 1.58.0 paused on 21-Jan-2024 while we investigate reports of a regression with portmapping.

All platforms

- The number of [4via6](#) site IDs are increased from 256 to 65,536
- [Taildrop](#) allows category Z unicode characters
- [DERP](#) flapping (flipping back and forth between two regions rapidly) is reduced when there's still an active connection for the home DERP server
- [Portmap](#) checks the epoch from NAT-PMP & PCP, and establishes a new portmapping if it changes
- Portmap better handles multiple interfaces
- Portmap handles multiple UPnP discovery responses
- Increased binary size with Tailscale 1.56 is resolved
- [Web interface](#) issue related to accessing shared devices
- Web interface login issue when accessed over HTTPS

Linux

- Shell shebang is added in postinstall script, which fixes some Debian installations

macOS

- **DNS Settings** view is added and displays the DNS configuration used when Tailscale is running
- Quit the app without terminating the VPN tunnel by holding down the **Option** button and selecting **Quit (Leave VPN Active)**
- **Toggle Tailscale** shortcut action can be used to connect or disconnect the VPN tunnel, depending on its current state
- The [KeyExpirationNotice](#) system policy is now supported to customize the time interval before a key expiration notice is displayed to the user
- The [web interface](#) is now supported in the [standalone](#) variant of the client
- Onboarding flow includes a step to ask the user to approve key expiry notifications
- Onboarding flow asks the user to approve the system extension if necessary, when using the standalone variant of the client
- Pre-Sonoma compatibility is improved
- VPN tunnel terminates upon closing the app
- Opening the **About** window triggers a check for updates
- The standalone variant of the client checks for updates every 72 hours

iOS

- **Toggle Tailscale** shortcut action can be used to connect or disconnect the VPN tunnel, depending on its current state. Ideal for the **Action Button** on iPhone 15 Pro.
- The [KeyExpirationNotice](#) system policy is now supported to customize the time interval before a key expiration notice is displayed to the user
- **Sign** button in the [Tailnet Lock](#) device sign view is rendered correctly
- Connectivity is no longer lost when transitioning from Wi-Fi to Cellular while an [exit node](#) is in use

Windows

- The [web interface](#) is now supported
- The lookup for net.sh.exe uses the absolute path instead of the relative path
- ADMX [system policy](#) descriptions are now available
- Vestigial wintun support is removed, which might have caused Chocolatey installs to break
- A goroutine leak in winMon no longer occurs if the monitor is never started
- "This package requires Windows 10 or newer" message no longer falsely displays during an uninstall or repair

Android

- Active network change detection is improved

tvOS

- Improvements to persistence of the client when running in the background

Kubernetes Operator

- A Connector custom resource is added, allowing users to configure the operator to deploy an [exit node](#), [subnet router](#), or both
- A warning displays if the unsupported ingress Exact path type is used
- StatefulSet labels are synced to their Pods
- A Tailscale [IngressClass](#) resource is added
- Extra long [Service](#) names are properly truncated

Containers

- Experimental support is added for configuring [tailscaled](#) using a mounted config file
- Tailscale images now contain layers of the same media type and can be parsed by [Podman](#) and [Buildah](#)

Jan 8, 2024

Tailscale v1.44.3

[Update instructions](#) →

Windows

- Added a security fix to address privilege escalation with [tailscale serve](#) and [tailscale funnel](#) that allowed low-privilege users to serve files they did not have access to (TS-2024-001). This release is intended for Windows 7 and 8 users. Those with later versions of Windows should run the latest stable version of Tailscale, which is 1.56.1. This issue was resolved in Tailscale 1.52.

Dec 15, 2023

Tailscale v1.56.1

[Update instructions](#) →

Linux

- [Web interface](#) redirects to the correct self-IP known by the source peer
- [App connector](#) domain list displays as expected

macOS

- [Custom login server](#) uses the provided URL instead of the Tailscale default login URL

iOS

- [Custom login server](#) uses the provided URL instead of the Tailscale default login URL

tvOS

- [Custom login server](#) uses the provided URL instead of the Tailscale default login URL

Dec 13, 2023

Tailscale v1.56.0

[Update instructions](#) →

All platforms

- `tailscale whois` command shows the machine and user associated with a Tailscale IP address
- [System policies](#) are now in [beta](#)
- `tailscale switch --list` command shows name and profile ID to disambiguate profiles with common login names
- Responsiveness is improved under load, especially with bidirectional traffic
- UPnP port mapping is improved

Linux

- The [web interface](#) allows users to configure some device settings such as [exit nodes](#), [subnet routers](#), and [Tailscale SSH](#) using a browser-based GUI instead of the [Tailscale CLI](#)
- `tailscale update` command is supported for [Unraid](#)
- `containerboot` symlinks its socket file if possible, making the Tailscale CLI work without `--socket=/tmp/tailscale.sock`

Windows

- Throughput is improved for [userspace \("netstack"\) mode](#) in the presence of packet loss
- Profile switcher displays the tailnet name
- Dynamic DNS updates are disabled in the client interface via the registry setting
- Client improvements when restarting after an upgrade

macOS

- [Taildrop](#) notification displays when a file is received (App Store variant only)
- Taildrop shortcut action is added for file sharing
- Profile switcher displays the tailnet name
- **About Tailscale** dialog indicates when the app is running a [TestFlight](#) build
- In-app warnings and push notifications display when internet connectivity is blocked because the current [exit node](#) is offline or its key has expired
- VPN tunnel fully terminates when Tailscale is stopped, using the menu bar toggle
- `/etc/resolv` file formatting with `Tailscaled-on-macOS` is improved

iOS

- **DNS Settings** view is added
- [Taildrop](#) shortcut action is added for file sharing
- Taildrop notifications include the received file names
- Profile switcher displays the tailnet name
- **About Tailscale** dialog indicates when the app is running a [TestFlight](#) build
- **Allow Local Network Access** option is added to the exit node picker UI

- In-app warning and push notification displays when internet connectivity is blocked because the current [exit node](#) is offline or its key has expired
- App size is reduced by about 2 MB with better asset compression

tvOS

- [Apple TV](#) can be configured as a subnet router, allowing you to remotely access resources on your home network that may not have Tailscale installed, such as a printer
- **About Tailscale** dialog indicates when the app is running a [TestFlight](#) build

Kubernetes

- [Helm charts](#) for the [Tailscale Kubernetes Operator](#) are now available on pkgs.tailscale.com/helmcharts
- [Kubernetes API server proxy](#) supports impersonating groups via [ACL Grants](#)
- Kubernetes operator [cluster egress](#) now supports referring to a tailnet service by its MagicDNS name in the Service annotation

GoKrazy

- TUN mode is used by default

Nov 30, 2023

Tailscale v1.54.1

[Update instructions](#) →

macOS

- Changing a pre-existing system policy value to nil no longer causes stability issues

iOS

- Changing a pre-existing system policy value to nil no longer causes stability issues
- Widget tracks the connection state more closely

tvOS

- Changing a pre-existing system policy value to nil no longer causes stability issues

Nov 15, 2023

Tailscale v1.54.0

[Update instructions](#) →

All platforms

- Go is updated to version 1.21.4

Linux

- [Substantially improve throughput](#) for UDP packets over TUN device with recent Linux kernels
- Added a security fix to address privilege escalation with `tailscale serve` and `tailscale funnel` that allowed low-privilege users to serve files they did not have access to if the machine administrator had previously granted that user `tailscale up --operator` privilege (TS-2024-001)

Windows

- Open menu with a regular select in addition to a right-select

macOS

- Implement MDM settings for the [standalone macOS application](#)
- Support for the `tailscale update` command for the standalone macOS application
- Don't run Taildrop cleanup loop until the first file transfer, and avoid spurious security dialog

iOS

- Show a helpful banner if there are no other devices on the tailnet
- Add **Allow Local Network Access** setting when using an exit node
- Show info bubble when key expires within 8 hours, or has expired
- Widgets now reflect the state of the VPN tunnel more accurately

QNAP

- Support for the `tailscale update` command

Nov 2, 2023

Tailscale v1.52.1

[Update instructions](#) →

Windows

- Resolve an incompatibility with other software that uses wintun

NAS platforms

- Clean up downloaded upgrades after applying them

Oct 30, 2023

Tailscale v1.52.0

[Update instructions](#) →**All platforms**

- `tailscale cert` command renews in the background. The current certificate only displays if it has expired.
- `tailscale status` command displays a message about client updates when newer versions are available
- `tailscale up` command displays a message about client updates when newer versions are available
- **Taildrop** now resumes file transfers after partial transfers are interrupted
- Taildrop prevents file duplication
- Taildrop detects conflicting file transfers and only proceeds with one transfer
- Wake on LAN (WoL) is now supported for peer node wake-ups
- TCP DNS queries are speculatively started if UDP hasn't responded quickly enough
- Truncated UDP DNS results are properly retried using TCP
- Go is updated to version 1.21.3

Linux

- `tailscale set` command flag `--auto-update` is added to opt in to automatic client updates (**beta**)
- `tailscale serve` and `tailscale funnel` commands are updated for improved usability
- `tailscale update` command for manual updates is now in **beta**
- **Taildrop** file transfer displays a progress meter
- `nftables` auto-detection is improved when `TS_DEBUG_FIREWALL_MODE=auto` is used
- DNS detection of NetworkManager with configured but absent `systemd-resolved`, such as EndeavourOS
- DNS detection for Debian `resolvconf` version 1.90 or later

Windows

- `tailscale set` command flag `--auto-update` is added to opt in to automatic client updates (**beta**)
- **Preferences** section contains auto-update setting
- Update notice displays, when a new version is available
- System policies allow system administrators to set a forced/suggested tailnet name, hide settings menu items, and more
- `tailscale serve` and `tailscale funnel` commands are updated for improved usability
- `tailscale update` command for manual updates is now in **beta**
- `iphlpsvc`, `netprofm`, and `winHttpAutoProxySvc` service dependencies are checked during installation
- Added a security fix to address privilege escalation with `tailscale serve` and `tailscale funnel` that allowed low-privilege users to serve files they did not have access to (**TS-2024-001**)

macOS

- `tailscale set` command flag `--auto-update` is added to opt in to automatic client updates (**beta**)
- App menu displays a notification item when a newer version is available
- System policies allow system administrators to set a forced/suggested tailnet name, prevent the VPN from stopping, hide categories of network devices and setting menu items, and more
- **Settings** section has an option added for turning on auto-updates
- **Reauthenticate** menu item shows time until expiry more prominently, presenting alerts when necessary
- `tailscale serve` and `tailscale funnel` commands are updated for improved usability
- `tailscale update` command for manual updates is now in **beta**

- **About** window more clearly distinguishes between the Standalone and App Store [variants](#) of the client
- Sparkle is updated to version 2.5.1

iOS

- **Settings** page displays a notification banner when a newer version is available on the App Store
- Home and lock screen widgets are supported
- System policies allow system administrators to set a forced/suggested tailnet name, prevent the VPN from stopping, hide the VPN On Demand settings, categories of network devices and settings menu items, and more

tvOS

- DNS support when operating as an [exit node](#)

Oct 2, 2023

Tailscale v1.50.1

[Update instructions](#) →

All platforms

- Fixed: `tailscale serve` configuration doesn't persist in container ([#9558](#))
- Tailnet Lock fails to sign node in container ([#9539](#))
- Funnel doesn't work for `tsnet` apps ([#9566](#))
- UPnP potentially crashes in specific circumstances

Sep 25, 2023

Tailscale v1.50.0

[Update instructions](#) →

All platforms

- [Wikimedia DNS](#) using DNS-over-HTTPS is supported
- Build with Go 1.21.1
- `tailscale update` command is unhidden on most platforms
- `tailscale ping` command sends an ICMP Ping code of 0
- `tailscale webcommand` updated to use React
- `tailscale debug portmap` command now has the `--log-http` option
- `tailscale netcheck` command works even if the OS platform lacks CA certificates
- UPnP falls back to a permanent lease if a limited lease fails
- WireGuard peer endpoint selections are improved

Linux

- Debian package lists the `iptables` and `iproute2` packages as recommended, not required

- [nftables](#) support interoperates with Uncomplicated Firewall (UFW)

Windows

- [tailscale bugreport](#) logs contain additional diagnostic information
- [Windows executable installer](#) detects when it is running on Windows 7 or Windows 8.x and will automatically download the appropriate v1.44.2 MSI package, which is the final release supporting those operating systems
- Windows executable installer no longer embeds MSI packages in the executable. Instead, it automatically downloads the correct package. Users desiring the previous behavior may download the "full" executable installer at pkgs.tailscale.com.

macOS

- Shortcuts are added for finding and pinging devices
- [Mullvad Exit Nodes](#) allows you to select nodes by country and city
- [Tailnet Lock](#) reliability improvements
- [Taldrop](#) no longer replaces spaces with %20 in file names when sending files to Windows devices

iOS

- [Fast user switching](#) is available
- iOS 17 supports customized device naming from **Settings**
- App Shortcuts in Spotlight and Siri are supported. Try saying: "*Hey Siri, connect to Tailscale*" or "*Hey Siri, is Tailscale connected?*".
- Shortcuts are added for finding and pinging devices
- Mullvad Exit Nodes includes an option to pick the best available node
- UI accessibility improvements when using VoiceOver
- Taldrop no longer replaces spaces with %20 in file names when sending files to Windows devices
- [VPN On Demand](#) rules are no longer reset when disabled and then restarted

Sep 11, 2023

Tailscale v1.48.2

[Update instructions](#) →

All platforms

- Stability improvements for [Mullvad Exit Nodes](#), particularly for users on IPv4-only networks

Aug 21, 2023

Tailscale v1.48.1

[Update instructions](#) →

All platforms

- Fix a security vulnerability in UPnP port mapping ([TS-2023-006](#))

Linux

- Fixed: Resolve nftables interaction between Tailscale and UFW which resulted in blocking subnet routed traffic

Synology

- Determine correct CPU architecture in `tailscale update` ([#8927](#))

Aug 16, 2023

Tailscale v1.48.0

[Update instructions](#) →

All platforms

- `tailscale exit-node` subcommand
- `--upstream` flag in the `tailscale version` command
- The `tailscale funnel` command provides an interactive web UI that prompts you to allow Tailscale to enable Tailscale Funnel on your behalf
- The `tailscale serve` command provides an interactive web UI that prompts you to allow Tailscale to enable HTTPS and Tailscale Funnel on your behalf
- [Tailnet Lock](#) is in [beta](#)

Linux

Note: 1.48.0 introduced a regression in the interaction between Tailscale and Linux `ufw`. The Linux release has been withdrawn pending a fix.

- Support for [nftables](#)
- RPM packages are now fully signed
- Support for the `tailscale update` command on Alpine, Arch and Fedora distro families

Synology

- Support for the `tailscale update` command

macOS

- Support for the `tailscale update` command

iOS

- Support for [VPN On Demand](#)
- VPN tunnel lifecycle improvements
- Improved exit node selection

- Minor UI tweaks

Jul 26, 2023

Tailscale v1.46.1

[Update instructions](#) →

All platforms

- Issue with [Tailnet Lock](#) signature verification

Linux

- Crash issue on ARM64

Android

- DNS and subnet routes issue

Jul 19, 2023

Tailscale v1.46.0

[Update instructions](#) →

Linux

- Initial support for nftables-based configuration. This option is currently behind a temporary flag for testing and feedback. See [issue #391](#) for details.

Windows

- [Tailnet Lock](#) is now supported

macOS

- Tailnet Lock is now supported

iOS

- Tailnet Lock is now supported
- Onboarding flow is added for easier initial setup of the app
- Ping devices on your tailnet from the app
- The app **Machines** page is improved
- The app **Exit Node** section is improved
- The app **Settings** page is improved

Jul 18, 2023

Tailscale v1.44.2

[Update instructions](#) →

All platforms

- Handling of custom HTTP ports in [tailscale serve](#)

Windows

- Restore support for Microsoft Windows 7 and Microsoft Windows 8.x.
Tailscale v1.44.2 will be the last release to support the following operating systems: Microsoft Windows 7, Microsoft Windows 8, Microsoft Windows Server 2008, and Microsoft Windows Server 2012.

Jul 10, 2023

Tailscale v1.44.1

[Update instructions](#) →

Android

- Various bugs and improvements

Jun 21, 2023

Tailscale v1.44.0

[Update instructions](#) →

Note: This is the last release to support the following operating systems:

- macOS 10.13 High Sierra
- macOS 10.14 Mojave

Tailscale releases after 1.44.0 will no longer install on these operating systems, though we expect to maintain forward compatibility and critical security updates for 1.44.0 with future releases until at least June 30, 2024.

To install Tailscale on a High Sierra or Mojave system, visit the Purchased Items in the App Store Account page. macOS High Sierra or Mojave systems will be offered Tailscale 1.44 when the download link is clicked. If Tailscale does not appear in the Purchased Items it must first be successfully installed using a recent macOS system. The Tailscale app will then be available for the High Sierra or Mojave system to install from Purchased Items.

All platforms

- [tailscale serve http](#) command to serve over HTTP (tailnet only)

- `tailscale ssh` command now supports remote port forwarding
- Recursive DNS resolution is now initially supported to replace bootstrapDNS when operating in a parallel mode
- Build with Go 1.20.5
- `--tun-userspace-networking` stability improvements for userspace subnet routers
- MagicSock private addresses are given preference when both private and public are available, to help keep traffic in private VPCs, where possible
- Async support is removed from the `portlist` package. Update to use synchronous `poll()` if this breaks your package.
- `watchIPNBus` now only requires read-only permissions to read
- `tailscale cert` renewal decision is now based on the lifetime of the certificate instead of hard-coded. This better supports 14 day certificate lifetimes.

Linux

- Changed:`tailscale ssh` support improvements for Security-Enhanced Linux (SELinux) systems
- Changed:`tailscale ssh` supports user names with up to 256 characters
- `build_dist.sh` better supports operating systems and CPU architectures which Tailscale release builds do not include
- The `iputils` package can now be installed on Alpine-based Docker containers

Windows

- PreferGo supports better DNS caching

macOS

- ICMP6 forwarding works as expected when running as a subnet router

FreeBSD

- ICMP6 forwarding works as expected when running as a subnet router

OpenBSD

- ICMP6 forwarding works as expected when running as a subnet router

WASI

- `tsnet` applications compiled to WebAssembly are now better supported

May 24, 2023

Tailscale v1.42.0

[Update instructions](#) →

Note: This is the last release to support the following operating systems:

- Microsoft Windows 7
- Microsoft Windows 8
- Microsoft Windows Server 2008
- Microsoft Windows Server 2012

Tailscale releases after 1.42.0 will no longer install on these operating systems, though we expect to maintain forward compatibility and critical security updates for 1.42.0 with future releases until at least May 31, 2024.

Note: Do not install this version of the Tailscale client on macOS 10.13. Upgrade to version 1.44.0 instead.

All platforms

- `tailscale serve reset` command to clear out the current serve configuration
- Changed: Update internal DNS handling to better support mixtures of global and private DNS servers

Linux

- SSH login on platforms which lack getent

Windows

Note: This release switches to a new application signing certificate, which is valid through 2025.

- Notification icons are updated

macOS

- Update Sparkle to check more regularly
- Taildrop delivery of incomplete files

iOS

- **Delete Account** button to redirect to the admin panel
- Better handle memory management to avoid hitting 50 MByte memory limit

Unraid

- Support Unraid as a NAS platform similar to how Synology and QNAP are handled

Kubernetes

- Support for `priorityClassName`

May 10, 2023

Tailscale v1.40.1

[Update instructions →](#)

Linux

- [Tailscale SSH](#) is now supported for LDAP users
- Support for Tailscale SSH session recording to a local file is restored
- Debian and RPM packages for MIPS architecture generate as expected

Windows

- Notification icons are updated
- The 32-bit [Windows installer](#) for the Tailscale client works as expected

macOS

- `tailscale cert` command no longer causes timeout failures

Kubernetes

- The Tailscale version displays in the startup logs

Apr 26, 2023

Tailscale v1.40.0

[Update instructions](#) →

All platforms

- `tailscale up --force-reauth` will now display a warning and 5 second countdown if you are connected over SSH over Tailscale, unless `--accept-risk=lose-ssh` is also given
- Tailscale now dynamically increases the buffer size for DERP relay messages based on the amount of available RAM ([#7776](#))
- Improvements were made to how Tailscale advertises available endpoints to reduce the likelihood of a spurious loss of direct connections ([#7877](#))

Linux

- Substantially higher throughput—for details, see [Surpassing 10Gb/s over Tailscale](#)
- Improved CPU consumption on systems with a very large (1M+) routing table

Windows

- Redo migration of pre-[Fast-User-Switching](#) state for better robustness

macOS

- "Settings" replaces "Preferences" as a menu item on macOS Ventura

Android

- Added intents `com.tailscale.ipn.CONNECT_VPN` and `com.tailscale.ipn.DISCONNECT_VPN`

gokrazy

- [Tailscale SSH](#) now works

QNAP

- UI failure after reboot

Apr 5, 2023

Tailscale v1.38.4

[Update instructions](#) →

All platforms

- Build with Go 1.20.3 to address security fixes ([CVE-2023-24537](#), [CVE-2023-24538](#), [CVE-2023-24534](#), and [CVE-2023-24536](#)). These address potential DoS attacks against DNS over HTTPS and Funnel that can occur over the public internet, and PeerAPI attacks launched from other nodes already on the tailnet.
- Added path support for proxy targets with `tailscale serve`
- Error displays when trying to use [Funnel](#) and `tailscale up --shields-up` simultaneously

Windows

- When connected to a Windows 10 client using [Windows RDP](#), the Tailscale taskbar right-click option for the remote client works as expected ([#7698](#))

Mar 29, 2023

Tailscale v1.38.3

[Update instructions](#) →

All platforms

- Support for stripping HTTP request paths from Funnel proxy routes ([#6571](#))
- [Tailscale Funnel](#) is now [beta](#)
- `tailscale serve` issue that did not use actual `SrcAddr` as `X-Forwarded-For`

Linux

- Certificate storage issue that did not actually use Kubernetes secrets

Windows

- Upgraded the Walk framework for the GUI client to improve menu responsiveness

Mar 22, 2023

Tailscale v1.38.2

[Update instructions](#) →

All platforms

- `tailscale lock tskey-wrap` has been replaced by `tailscale lock sign`
- `tailscale lock sign` now supports signing auth keys

Linux

- `--tun=userspace-networking` issue running in [Azure App Services](#)

macOS

- Sparkle automatically checks [updates for the standalone package](#). This does not impact the App Store package.

FreeBSD

- Fixed: Issue setting the effective group ID on some non-interactive Tailscale SSH sessions. This issue is specific to FreeBSD's implementation of `setgroups` and does not impact other platforms.

Mar 14, 2023

Tailscale v1.38.1

[Update instructions](#) →

All platforms

- `tailscale configure` command to configure resources that you want to include in your tailnet
- `tailscale lock sign` to sign [pre-approved auth keys](#) for use with [Tailnet Lock](#)
- `tailscale debug derp` command to help diagnose DERP-related difficulty
- `tailscale debug capture` command to write packet capturing for debugging
- The `tailscale debug portmap` command replaces `tailscaled debug -portmap`. This is now available on platforms without a `tailscaled` binary (like the macOS App Store).
- `tailscale serve` command has been overhauled
- `tailscale serve funnel` has been made into its own command, [tailscale funnel](#)
- Several improvements to UPnP port mapping have been made that allow it to work with a broader set of home routers

Linux

- Certificates can be stored in Kubernetes secret storage

Windows

- MSI installers start the GUI without user interaction to allow remote upgrades

macOS

- Notification upon node key expiration (only on macOS 10.14 and later)
- [Tailscale SSH](#) server component is available for macOS open source [Tailscale + tailscaled CLI devices](#)

iOS

- Support for alternate control servers by setting the URL in Settings page of the admin console

Android

- Chromecast support while Tailscale is active

Note: v1.38.0 was never released.

Feb 22, 2023

Tailscale v1.36.2

[Update instructions](#) →

macOS

- Prevent using an exit node while being an exit node
- Improve detection of default interface

iOS

- Improve detection of default interface

Windows

- Improve clean out of registry entries during upgrade

Feb 8, 2023

Tailscale v1.36.1

[Update instructions](#) →

All Platforms

- Potential infinite loop when node key expires

macOS

- Handle starting the app before network interfaces are ready

iOS

- Handle starting the app before network interfaces are ready
- Get Status intent will not connect the VPN

Windows

- Potential crash in netstat handling
- Windows 7 checks for KB2533623

Jan 24, 2023

Tailscale v1.36

[Update instructions](#) →

All Platforms

- `--json` flag for the `tailscale lock status` and `tailscale lock log` commands
- `--json` flag for the `tailscale version` command
- `tailscale update` command to update client
- `tailscale debug daemon-logs` to watch server logs
- `tailscale status --json` now includes `KeyExpiry` time and `Expired` boolean on nodes
- `tailscale version` now advertises when you're on the `unstable` (dev) track
- (Unix platforms) When `/etc/resolv.conf` needs to be overwritten for lack of options, a comment in the file now links to tailscale.com/s/resolvconf-overwrite
- **Tailscale SSH**: SSH to tailscaled as a non-root user works again, as long as you only SSH to the same user that tailscaled is running as
- Fixed: Handle cases where a node expires and we don't receive an update about it from the control server ([#6929](#) and [#6937](#))
- Fixed: Support UPnP port mapping of gateway devices where they are deployed as a highly available pair ([#6946](#))
- Support arbitrary IP protocols like EOIP and GRE ([#6423](#))
- Exit node handling of a large number of split DNS domains ([#6875](#))
- Accept DNS-over-TCP responses up to 4K bytes ([#6805](#))

Linux

- Add build support for Loongnix CPU architecture
- **Improved throughput performance** on Linux ([#6663](#))

macOS

- Tailscale actions (connect, disconnect, switch profile, use exit node) are available in the Shortcuts app (read the [blog post](#))
- Tailscale traffic looping upon certain sleep/resume/Wi-Fi change transitions ([#5156](#))

iOS

- Tailscale actions (connect, disconnect, use exit node) are available in the Shortcuts app
- Tailscale using cellular data even after Wi-Fi becomes available ([#6565](#))

Windows

- Add a more robust mechanism to remove WinTun ([#6433](#))
- Update taskbar menu radio button implementation

Android

- New version of the Gio UI library with internationalization and accessibility fixes
- Allow Sonos app to discover local devices while Tailscale is connected

Synology

- Show whether outgoing connections are configured in the web UI

Containers

- Run in a Kubernetes environment without setting TS_KUBE_SECRET ([#6704](#))

OpenBSD

- [Tailscale SSH](#) runs on OpenBSD

Jan 4, 2023

Tailscale v1.34.2

[Update instructions](#) →

Linux

- Handling of a very large number of SplitDNS domains with an exit node

macOS

- UI glitch with macOS 10.14 and 10.13

Windows

- Custom server URL from registry key support

Synology

- Crashes manifesting on ARM-based platforms and models with very old kernels

Dec 13, 2022

Tailscale v1.34.1

[Update instructions](#) →

Linux

- Unit tests on systems using busybox ip
- Regression handling TS_STATE_DIR in containerboot

macOS

- Issue which could fail to save the key for `tailscale serve` (#6409)
- Issue which could cause crash when interfaces change (#6641)

Windows

- Common cause of an issue with [Tailscale SSH](#) (#6639)

Dec 5, 2022

Tailscale v1.34

[Update instructions](#) →

All Platforms

- `tailscale switch` command to switch between accounts using [fast user switching](#)
- `tailscale login` command to [login with a specified account](#)
- `tailscale set` command to modify configuration settings without needing to repeat the others
- `tailscale lock` command to manage [Tailnet Lock](#) for your tailnet
- Additional [4via6 DNS name](#) format, Q-R-S-T-via-X (or Q-R-S-T-via-X.yak-bebop.ts.net), for systems that required dashes instead of dots
- Display decoded punycode hostnames in status list
- Warn in `tailscale status` health and `tailscale up` if there are nodes advertising routes but --accept-routes=false

Linux

- Add [fast user switching](#) using `tailscale login` and `tailscale switch`
- Warn in `tailscale status` health if something else overwrites /etc/resolv.conf

macOS

- Add [fast user switching](#) by selecting the desired tailnet from the Tailscale icon in the menubar, or via the `tailscale login` and `tailscale switch` commands

Windows

- Add [fast user switching](#) by selecting the desired tailnet from the Tailscale icon in the taskbar, or via the `tailscale login` and `tailscale switch` commands
- Use named pipes to communicate between UI and Service
- Changed: Move state storage responsibility from frontend to backend. The current state is migrated, this should not be a noticeable change.
- Switch to wingoes for OLE support, use multithreaded apartment
- Received [Taildrop](#) files get placed in the `C:\Users\\Downloads` directory (previously they were placed in the `C:\Users\\Desktop` directory)

Android

- Allow Sonos app to discover speakers on the local LAN

Synology

- Better detect DSM version, locate local socket correctly

Containers

- Replace `run.sh` with `cmd/containerboot`

FreeBSD

- Support for [Tailscale SSH](#) (Thanks Pat Maddox!)

Nov 21, 2022

Tailscale v1.32.3

[Update instructions](#) →

All Platforms

- Security vulnerability in the Windows client that allows a malicious website to reconfigure the Tailscale daemon `tailscaled`, which can then be used to remotely execute code ([CVE-2022-41924](#), [TS-2022-004](#))
- Fixed: Security vulnerability in the client that allows a malicious website to access the peer API, which can then be used to access Tailscale environment variables ([CVE-2022-41925](#), [TS-2022-005](#))

Windows

- Set `Zone.Identifier` alternate data stream for Taildrop files

macOS

- Set `com.apple.quarantine` flag for Taildrop files

Oct 26, 2022

Tailscale v1.32.2

[Update instructions](#) →

All Platforms

- Substantially improve userspace-networking handling of packet loss

macOS

- Fix a crash impacting some macOS systems ([#6065](#))

Android

- Fix a 4-in-6 DNS problem mainly impacting Android (fixed by Peter Cai) ([#5698](#))

Oct 21, 2022

Tailscale v1.32.1

[Update instructions](#) →

All Platforms

- Avoid crash in tailscale netcheck ([#5919](#))

macOS

- Avoid a condition which can result in high CPU consumption ([#5879](#))
- Fix Taildrop failures when sending many files ([#5873](#))

iOS

- Fix Taildrop failures when sending many files ([#5873](#))

Windows

- Correct IPv6 MTU setting ([#5914](#))

Oct 12, 2022

Tailscale v1.32.0

[Update instructions](#) →

All Platforms

- Support NextDNS
- Add tailscaled --no-logs-no-support (or TS_NO_LOGS_NO_SUPPORT=true environment variable)
- tailscale bugreport --record flag to pause and write another bug report
- More in-depth health checks in a bugreport
- tailscale netcheck looks for a captive portal
- Build with Go 1.19.2
- IP fragmentation handling as an exit node
- SSH inadvertently closing tmux/etc panes at disconnect
- Always respond to 4via6 ICMP echo requests
- Normalize more process names in Services report

Linux

- Coexist with mwan3 package iptables rule fwmark masks, for OpenWRT
- Add an eBPF helper to pass the first packet on a new flow up to tailscaled
- Better detect when running in a container

macOS

- Incorrect list of Taildrop target devices

Windows

- Log Windows service diagnostics when the wintun device fails to install

iOS

- Incorrect list of Taildrop target devices

Android

- Show an error when unable to accommodate multiple users

Synology

- envknob support
- Configure-host version parsing

Sep 16, 2022

Tailscale v1.30.2

[Update instructions](#) →

All Platforms

- IPv6-mapped-IPv4 addresses in STUN responses

- Better detect when running in a container

Sep 8, 2022

Tailscale v1.30.1

[Update instructions](#) →

All Platforms

- Exit nodes in userspace-networking mode break Chrome v.104 or later IPv6 connectivity
- SIGINT when running in a container without job control

Aug 31, 2022

Tailscale v1.30

[Update instructions](#) →

All Platforms

- Use DNS-over-HTTPS for Mullvad DNS servers
- Report whether a subnet router is running in userspace-networking or kernel mode
- send Tailscale client version number in ACME requests (to Let's Encrypt, for example)
- Report whether host kernel supports IPv6
- Add `tailscale` licenses with link to open source licenses
- Delete node immediately if `tailscaled` exists and was using `mem: state storage`
- `tsnet` ephemeral nodes will delete themselves on `close()`
- Add a timeout when writing to BIRD socket
- Clients can use Noise with any HTTPS port with `capver 39` (mainly for Headscale)
- `100.100.100.100` will respond with `SERVFAIL` if there are no upstream resolvers

Linux

- Gracefully handle restarts in resolved support

macOS

- Report `variant` (App Store, system extension) in the about box
- Fix missing IP address display in the status menu

Windows

- Add native ARM build for backend Tailscale service (only in NSIS installer in this release)
- Update Proxy support
- Notice when group policy entries change and move our NRPT rules between the local and group policy subkeys as needed
- Avoid 2.3 second DNS lookup delay when Smart Name Resolution is enabled by adding MagicDNS names to hosts file

- Disable NetBIOS nameservice on Tailscale interfaces

iOS

- Fix potential crash in notification handling
- Fix dismissing of error indication if a bugreport fails

Android

- Allow coordination server URL to be set. Select the Authentication menu three times quickly to enable
- Fix Google Stadia, Android Auto, GoPro, and Messages RCS with the VPN active

Synology

- Fix `/dev/net` permissions in `tailscale configure-host`

OpenBSD

- Support functioning as a subnet router or exit node using hybrid netstack mode

Other

- Accommodate shared nodes in `nginx-auth`
- Fix race in `derper` (Custom DERP servers) with manual certificates

Jul 18, 2022

Tailscale v1.28

[Update instructions](#) →

All Platforms

- Add `ExitNodeStatus` to `tailscale status --json`
- Fix `tailscale ping -c N` to properly exit after `N` ping requests even if there are timeouts
- MagicDNS recursive resolution now returns `SERVFAIL` if all upstream resolvers fail
- `portmapper`: Send discovery packet for IGD specifically, some routers don't respond to `ssdp:all`

Linux

- Implement specific DNS support for AWS, Google Cloud, and Azure to add internal split DNS domain and fallback DNS

macOS

- Use one large `100.64.0.0/10` route entry if there are no other interfaces using CGNAT, to avoid Network Changed errors in browsers where possible

Windows

- Suppress nonfunctional link-local IPv6 addresses on Tailscale interface, PowerShell ping (hostname) now works correctly
- Set registry values to not send DNS changes concerning our interface to AD domain controllers
- Update Windows split DNS settings to work alongside other NRPT entries set by group policy
- Set AllowSameVersionUpgrades attribute on MajorUpgrade tag in Windows MSI script

iOS

- Add portmapper support for NAT-PMP, PCP, UPnP
- Add MagicDNS support for TCP
- Changed: The minimum iOS version is now iOS 15, which makes substantially more memory available (the App Store will offer Tailscale 1.26.2 for iOS 13 and 14 devices)

Android

- Android can now be an exit node (previously available but hidden)

Jul 5, 2022

Tailscale v1.26.2

[Update instructions](#) →

All Platforms

- tailscaled being able to restart while mosh-server is running from an SSH session
- Make tailscale up --operator="" clear a previously set operator

Linux

- [Tailscale SSH](#) support with Arch Linux

macOS

- Limit SSH login to 16 groups

Windows

- Make SSH command prefer Windows ssh.exe over PATH

iOS

- Try harder to notify for [SSH check mode](#)

Jun 18, 2022

Tailscale v1.26.1

[Update instructions](#) →

All Platforms

- Various bugs

Jun 6, 2022

Tailscale v1.26

[Update instructions](#) →

All Platforms

- Add `--peerapi <peer>` flag in `tailscale ping` to check connectivity to a peer using the PeerAPI
- Add `--timeout <duration>` flag in `tailscale up` to enforce a maximum amount of time to wait for the Tailscale service to initialize
- Allow `LoginInteractive` via `LocalAPI`
- MagicDNS supports DNS/TCP and handling IP fragmented UDP frames
- Add an overall 10 second timeout for recursive MagicDNS queries
- Add `wake-on-LAN` function to PeerAPI. There is no UI for it currently.
- Provide `/run.sh` as an entrypoint for Docker container builds
- Configured MTU is now consistent between a TUN device and a userspace device
- Refactor `tailscale.com/client/tailscale` package with `LocalClient` type
- Change MagicDNS "via route" DNS names from "via-SITEID.10.2.3.4" to "10.2.3.4.via-SITEID". The old format will continue to work for the next one or two releases.
- Build with Go 1.18.3

macOS

- [Tailscaled-on-macOS](#) now supports MagicDNS, including Split DNS
- Initial release of a standalone macOS client, which is independent of the App Store, in the [stable track](#)

Windows

- Add `TS_NOLAUNCH` property to allow admins to deploy silent MSI installs without automatically starting the GUI
- MagicDNS lookup of own hostname
- Handle more than 50 Split DNS domains
- Resolve one source of shutdown delay (there may still be more)

Synology

- Allow the NAS disks to hibernate by moving telemetry buffering to `tmpfs`
- Improve HTTP proxy handling

iOS

- **Bug report** menu option in the UI

Apr 28, 2022

Tailscale v1.24.2

[Update instructions](#) →

All Platforms

- Handling of HTTP proxies in certain circumstances
- An issue where the new control plane protocol could fail to make a connection to our servers ([#4557](#))

Synology

- Additional fix in handling of HTTP proxies

Apr 27, 2022

Tailscale v1.24.1

[Update instructions](#) →

All Platforms

- Two issues where the new control plane protocol could fail to make a connection to our servers ([#4544](#), [#4538](#))
- Set TCP keep-alives in userspace-networking subnet router to avoid connection leaks ([#4522](#))
- Avoid using the LTE radio after transition to Wi-Fi

Apr 22, 2022

Tailscale v1.24

[Update instructions](#) →

All Platforms

- Initial support for site-relative IPv4 addressing using IPv6
- First for-keepsies deployment of ts2021 protocol
- tsnets now supports providing a custom ipn.StateStore
- Improve netstack performance via better GC tuning
- MagicDNS: PTR records for TS service IPs
- Build with Go 1.18

Linux

- `taildrop: add file get --loop`

- taildrop: add file get --conflict=(skip|overwrite|rename)
- Default to userspace-networking mode on [gokrazy](#)
- Set tailscale0 link speed to UNKNOWN, not 1Gbps
- Attempt to load the xt_mark kernel module when it is not present

Windows

- Improve HTTPS proxy handling

Synology

- Improve HTTPS proxy handling

Android

- Android TV support
- Fix and reintroduce Talkback support

FreeBSD

- Portmapping support

Mar 17, 2022

Tailscale v1.22.2

[Update instructions](#) →

Linux

- Potential crash at startup when using BGP

Windows

- MSI not restarting GUI after MSI-to-MSI upgrade

Mar 9, 2022

Tailscale v1.22.1

[Update instructions](#) →

All Platforms

- In userspace-networking mode, always close SOCKS proxied connections

Linux

- Better operation with [gokrazy](#)

macOS

- Fix macOS GUI "Must restart" dialog in some cases

Windows

- Fix a Windows NSIS installer bug when upgrading

FreeBSD

- Fix portmapping

Feb 23, 2022

Tailscale v1.22

[Update instructions](#) →

All Platforms

- New: DERP Return Path Optimization (DRPO), allows a pair of nodes in different DERP regions to connect more quickly by only requiring one side to connect to the other, cutting down some DERP setup latency
- `tailscaled --state=mem`: registers as an ephemeral node and does not store state to disk
- `tailscale status --json` now shows `Tags` and `PrimaryRoutes` for Peers. `PrimaryRoutes` shows whether a HA subnet router is currently the active one.
- `tailscale status --json | jq .TailnetName` will show the name of the tailnet
- The optional `tailscaled` debug server's Prometheus metrics exporter now also includes Go runtime metrics
- `tailscaled` supports a new `TS_PERMIT_CERT_UID` environment variable containing either a `userid` or `username` to allow to fetch Tailscale TLS certificates for the node. This environment variable can be set in `/etc/default/tailscaled` to permit non-root web servers on the local machine to fetch certs from `tailscaled`.
- Send heartbeats less often, saving some battery, matching v1.20 change on mobile platforms.
- `--auth-key` and `--authkey` both work as `tailscale up` arguments

Linux

- More robust detection of `systemd-resolved`
- Efficiently parse extremely large `/proc/net/route` files
- Be more helpful in suggesting `tailscale --operator=USER` to use with Taildrop
- Some broken host DNS configurations are now detected and reported in `tailscale status`

Windows

- MSI installer

- Reject SIDs from deleted/invalid security principals to avoid failed to look up user from userid error

Synology

- Add `/var/packages/Tailscale/target/bin/tailscale configure-host` to restore needed permissions. We recommend adding this as a scheduled task at boot.

Feb 9, 2022

Tailscale v1.20.4

[Update instructions](#) →

All Platforms

- DNS lookups via an exit node in many cases

Linux

- Better handling of extremely large `/proc/net/route` files for very large routers
- BGP advertisement with subnet router failover

OpenBSD

- `openresolv /etc/resolv.conf` handling

Jan 26, 2022

Tailscale v1.20.3

[Update instructions](#) →

Synology

- UI issues in Synology ([Synology 1.20.2 doesn't have working options page](#))

Only the Synology client released v1.20.3. All other platforms remain with v1.20.2.

Jan 21, 2022

Tailscale v1.20.2

[Update instructions](#) →

All Platforms

- Memory footprint growth in userspace-networking mode ([netstack: leaking packet buffers tailscale #3762](#))
- Userspace-networking will accept a TCP SYN with ECN bits set ([xt-userspace-networking incoming TCP doesn't always work right away tailscale #2642](#))
- Saving resolver list for OpenBSD

Jan 13, 2022

Tailscale v1.20.1

[Update instructions](#) →

All Platforms

- Deadlock in handling the DERP map

Jan 12, 2022

Tailscale v1.20

[Update instructions](#) →

All Platforms

- When using an exit node, DNS queries will be forwarded to the exit node for resolution
- `tailscaled` now allows running the outgoing SOCKS5 and HTTP proxies on the same port.
- SOCKS5/HTTP proxies now allow connecting via subnet routers & exit nodes when run in userspace-networking mode
- More debug metrics available
- `tailscale ip -1` flag
- CLI now lets you select exit node by name
- CLI now shows you which nodes are offering exit nodes
- CLI now refuses to let you pick an invalid exit node (when connected)
- Packet filter now supports matching any IP protocol number when enabled in ACLs (previously only TCP, UDP, ICMP and SCTP)
- Added `online` boolean to `tailscale status --json`, made `tailscale status` show offline nodes
- Added `tailscale up --json`
- MagicDNS now works over IPv6 when CGNAT IPv4 is disabled using `disableIPv4: true` in ACL
- Choose a new DERP relay server if the current DERP is removed from the DERPmap
- Bug fixes, cleanups, log spam reduction

Linux

- `tailscale file cp` sends via the local `tailscaled` now, so it now supports `tailscaled` running in tun-free, userspace-networking mode (such as on Synology DSM7 unless you [enable TUN mode](#))

Windows

- GUI support for running an exit node

macOS

- GUI support for running an exit node

iOS

- Send heartbeats less often to conserve battery

Android

- Talkback support
- Menu selection to generate a bug report
- "Allow LAN Access" checkbox in Exit Node menu
- Send heartbeats less often to conserve battery
- Implement DNS config reporting
- No longer require fallback DNS to be configured in admin console
- Report in the UI when connectivity is lost; this functionality was present but broken in prior releases

FreeBSD

- Now supports running in a jail (if devd isn't available, it falls back to network status polling mode)

Dec 15, 2021

Tailscale v1.18.2

[Update instructions](#) →

All Platforms

- Permit protocols other than TCP, UDP, or SCTP if an [ACL rule](#) has a proto specified and allows * port range
- Exit node selection takes effect (almost) immediately

Linux

- In DNS DirectManager, allow comments at the end of a line
- Don't get stuck waiting for systemd-resolved to restart in one particular DNS configuration

Synology

- [Receive Taildrop files](#)

Nov 25, 2021

Tailscale v1.18.1

[Update instructions](#) →

Linux

- Regressions on some kernel configs related to our direct use of netlink rather than using the `ip` command to program routes and policy routing

Nov 18, 2021

Tailscale v1.18

[Update instructions](#) →

All Platforms

- `tailscaled` debug server now exports Prometheus metrics at `/debug/metrics`
- Improved UPnP discovery so that eero devices now work, allowing a port to be opened for direct connections (also in 1.16.2)
- State machine transition regarding expired key extension
- If unable to upload telemetry, limit amount buffered to 50MB
- Retry more transient DNS errors, instead of passing the failure back to the client

Linux

- Support storing Tailscale state using AWS SSM (for example, `tailscaled -state arn:aws:ssm:eu-west-1:123456789:parameter/foo`) (thank you Maxime Visonneau)
- If `resolvconf` wrote `/etc/resolv.conf` but pointed it to `systemd-resolved`, use `systemd-resolved` for DNS not `resolvconf`
- If `NetworkManager` wrote `/etc/resolv.conf` but pointed it to `systemd-resolved`, use `systemd-resolved` for DNS not `NetworkManager`
- Handle `/etc/resolv.conf` being a bind mount into a container, such that we cannot `rename()` it.
- Work around Ubuntu 18.04 `setLinkDomain` length limit by omitting reverse lookup information
- Use `AF_NETLINK` messages to configure IP, not the `ip` command. Set `TS_DEBUG_USE_IP_COMMAND` environment variable to revert to use of `/sbin/ip` if this breaks your device

iOS

- On iOS 15+, where Network Extensions have more memory available, allow the same number of DNS-over-HTTPS requests in flight as other platforms

Synology

- Only use `AmbientCaps` on DSM7+

Oct 7, 2021

Tailscale v1.16

[Update instructions](#) →

All Platforms

- Support storage of node state as a Kubernetes secret.
- `tailscale up --authkey=file:/path/to/secret` support
- `tailscale up --qr` for QR codes
- `tailscaled` in userspace-networking mode can now run an HTTP proxy server (in addition to the prior SOCKS5 proxy server support)
- No longer need the `while tailscale up; do sleep 0.1; done` loops in Docker startup scripts.
- CPU/memory profiling support in `tailscale debug`
- Bake in Let's Encrypt's ISRG Root X1 root (also in 1.14.6)

Linux

- Support containers with `!CAP_NET_RAW` and `!CAP_NET_ADMIN` (like CircleCI runners)
- Service (portlist) scanning optimized; uses much less CPU on busy servers

Windows

- Move state to `C:\ProgramData` (also in 1.14.4)

macOS

- Super rare Wireguard packet loop network flood when using a DNS server behind a subnet router, when a macOS device resumes from sleep and the network changes (also iOS, but triggers less there). Fixes [#1526](#) (also in 1.14.6).

iOS

- Turn the radio on less often to improve battery performance

Android

- Support Taildrop on older Android releases
- Turn the radio on less often to improve battery performance

Oct 1, 2021

Service Updates

[Update instructions](#) →

All Platforms

- Include Let's Encrypt's ISRG Root X1 root as an alternate to try if the platform roots fail
- If `tailscale cert` fails because it needs to be run as root, say so.
- Avoid looping packets in `tstun`, believed to fix [#1526](#)
- Allow SOCKS5 proxy for `--tun=userspace-networking` to dial the HTTPS domain name of the Tailnet
- Ensure state directory is set to perm 0700.

iOS

- Ignore ipsec link monitor events for iOS to avoid waking the system

Sep 24, 2021

Service Updates

[Update instructions](#) →

Windows

- Move state files from C:\Windows to C:\ProgramData, to better handle Windows

Synology

- Fix segfaults shortly after starting, resolves [#2733](#)

Sep 16, 2021

Tailscale v1.14.3

[Update instructions](#) →

All Platforms

- `tailscale up` will wait for the socket to `tailscaled` to be created, not exit with an error. It should no longer be necessary to run it in a loop.
- Crash in TCP forwarding with `userspace-networking`; resolves [#2658](#)

Windows

- Default route lookup on Windows; resolves [#2707](#)

Note: v1.14.1 and v1.14.2 were never released.

← [rss-offline](#) · [abrir original](#) · **OpenClaw Is Getting Faster, Smaller, and Easier to Trust**
[rendered] · OpenClaw Blog · 2026-05-27 21:00



OpenClaw

[Product](#)

[Ecosystem](#)

[Docs](#)

[Blog](#)

[GitHub](#)

All posts

OPENCLAW BLOG

OpenClaw Is Getting Faster, Smaller, and Easier to Trust

A release sweep across February through May shows faster agent turns, fewer dependencies, and a cleaner package shape.



Peter Steinberger

[@steipete](#)


May 28, 2026 4 min read

OpenClaw has been getting faster and smaller at the same time. The performance work is visible in agent turns. The dependency work is quieter, but it cuts npm size, install size, audit surface, and native package surprises.

The package grew while OpenClaw gained channels, providers, media, memory, and plugin SDK surface. Then we started moving heavier plugin dependency cones out of core. The full release rows and caveats live in the [technical report](#).

STABLE COLD TURN

2.9x faster



V2026.4.14


V2026.5.27

9.8S

3.4S

STABLE WARM TURN

2.5x faster



V2026.4.14


V2026.5.27

7.5S

3.0S

AGENT PEAK RSS

7% lower



V2026.4.14


V2026.5.27

686 MB

635 MB

PUBLISHED TARBALL

59% smaller






2026.3.31

2026.5.27

43.3 MB

17.8 MB

Cold agent turn trend



335.0 MB

5.22*

1,020.6 MB

5.27*

786.9 MB

* Root shrinkwrap landed in 5.22; the size jump came from a bad package shape that made npm install a duplicate dependency tree.

Installed dependency count

2.26

645

3.31

438

5.22

401

5.27

371

Already down to 314 on main for the next release.

INSTALLED DEPENDENCIES

371

Latest release, down 42% from the monthly high.

Main is already at 314.

DUPLICATE INSTALL COPY

found

5.27 still shows the shrinkwrap-exposed duplicate tree.

Removed on main for the next release.

SHRINKWRAP

stays

Shrinkwrap was not the problem; the package shape was.

February and March

More product, larger package

The npm package grew from 82.9 MB unpacked to 182.6 MB unpacked while the surface area expanded.

2026.5.12

Plugin extraction becomes visible

Bedrock, Slack, OpenShell, Anthropic Vertex, Matrix, and WhatsApp move out of the core dependency path.

2026.5.22

Shrinkwrap exposed bad package shape

npm materialized a large nested tree with every canvas platform package.

2026.5.27

Latest release: smaller package, known install debt

17.8 MB published tarball, 371 installed dependencies, and the shrinkwrap-exposed duplicate tree still visible in fresh installs.

Already removed on main for the next release.

The direction is simple: keep core small, move optional capabilities into plugins, make dependency ownership explicit, and measure the user-visible effects. Each point is one smoke run, useful for spotting large shifts rather than making fine-grained benchmark claims.

For methodology, caveats, per-release rows, and the shrinkwrap boundary audit, read the [technical report](#).

Growth, here, looks more like molting than adding.

Share this post



Hacker News

Keep reading

May 5, 2026

[OpenClaw Had a Rough Week](#)

May 15, 2026

[Where OpenClaw Security Is Heading](#)

Apr 30, 2026

[How OpenClaw Got Safer in Public](#)


Build from the docs


Install, configure, and harden OpenClaw with the same reading layout.

[Open docs](#)




[Inside the Architecture of DigitalOcean's Inference Router](#)






[Now Available: Kimi K2.6](#)



[Missed the Deploy 2026 Keynote live? Watch it now and catch everything that was announced](#)



- [Blog](#)
- [Docs](#)
- [Careers](#)
- [Get Support](#)
- [Contact Sales](#)



- [Products](#) 

Featured AI Products

>

Compute

Build, deploy, and scale cloud compute resources

>

Containers and Images

Safely store and manage containers and backups

>

Managed Databases

Fully managed resources running popular database engines

>

Management and Dev Tools

Control infrastructure and gather insights

>

Networking

Secure and control traffic to apps

>

Security

Help protect your account and resources with these security features

>

Storage

Store and access any amount of data reliably in the cloud

>

[Browse all products](#)

•

AI/ML

>

CMS

>

Data and IoT

>

Developer Tools

>

Gaming and Media

>

GPU

>

Hosting

>

Security and Networking

>

Startups and SMBs

>

Web and App Platforms

>

[See all solutions](#)

•

Community

>

Documentation

>

Developer Tools

>

Get Involved

>

Utilities and Help

>

• [Partners](#) ▾

Become a Partner

>

Marketplace

>

• [Pricing](#)

• [Log in](#)
• [Sign up](#)

• [Log in](#)
• [Sign up](#)



[Product updates](#)

Scalable, Cost-Efficient AI: Introducing Unified Batch Inference on DigitalOcean



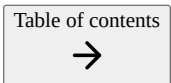
By [snamdeo](#) and [smirza](#)

- Updated: May 28, 2026
- 8 min read

[<- Back to blog home](#)

Table of contents

1. [The AI Scaling Bottleneck](#)
2. [Introducing DigitalOcean Batch Inference](#)
3. [Deeply Integrated with DigitalOcean](#)
4. [How It Works](#)
5. [Use Cases](#)
6. [Getting Started](#)
7. [The Bigger Picture](#)



[The AI Scaling Bottleneck](#)[Introducing DigitalOcean Batch Inference](#)[Deeply Integrated with DigitalOcean](#)[How It Works](#)[Use Cases](#)[Getting Started](#)[The Bigger Picture](#)

At Deploy 2026, we introduced the DigitalOcean AI-Native Cloud, built for the inference era. [Batch Inference](#) on the [DigitalOcean Inference Engine](#) enables high-volume asynchronous workloads. As developers move from AI prototypes to production-scale applications, the challenges of cost and rate limits often become a bottleneck. Batch Inference addresses these hurdles by allowing you to process high-volume workloads asynchronously at a fraction of the cost of synchronous requests.

Whether you are performing large-scale data transformation, content generation, building embeddings or offline evaluations, Batch Inference provides a unified, consistent way to leverage the world’s leading models from OpenAI and Anthropic, all through a single DigitalOcean interface.

The AI Scaling Bottleneck

Real-time inference is essential for interactive AI applications such as chatbots, copilots, and search-as-you-type experiences. However, when the task involves processing 10,000 support tickets for sentiment analysis, generating SEO metadata for an entire product catalog, or benchmarking a new system prompt against a test suite, real-time inference becomes an expensive and inefficient tool for the job.

Each of those requests competes for the same rate-limited throughput as your production traffic. Teams spend engineering time writing retry logic, managing backpressure, and monitoring scripts that work through sequential API calls for hours. If you use models from multiple providers, such as OpenAI for embeddings and Anthropic for generation, you are managing separate credentials, separate billing dashboards, and separate error-handling strategies, even though the core workflow is the same: submit requests, wait, retrieve results.

Processing thousands of synchronous requests is not only slow, it is an architectural challenge. At scale, synchronous inference becomes inefficient requiring thousands of open connections, creating constant rate-limit pressure and wasting compute while waiting for responses. It also introduces throughput bottlenecks, retry storms, and inconsistent latency, while pushing complex orchestration logic (queuing, retries, backoff) onto the client. Across multiple providers, this fragmentation only compounds the operational burden.

Introducing DigitalOcean Batch Inference

With Batch Inference, you can submit up to [50k requests for OpenAI](#) or [100k for Anthropic](#) in a single `.jsonl` file and let DigitalOcean handle the orchestration: queuing, execution, and result delivery.

What distinguishes this approach is its unified interface. **Instead of working with each provider individually, OpenAI and Anthropic models are accessible through a single DigitalOcean API.** One set of endpoints, one authentication flow, and one billing account allows you to monitor every job in one place, regardless of which provider executes it.

This single control plane manages the operational complexity while preserving full access to each provider's native model capabilities.

DigitalOcean Batch Inference provides a single control plane

The upload, submission, and retrieval workflow is identical regardless of which model you use. By using one set of endpoints and a single authentication flow, you can switch between (or combine) providers without rewriting your orchestration logic or reconciling separate invoices.

Significant Cost Savings

Batch requests are billed at a significant discount compared to standard real-time inference rates, for both input, output, and cache tokens. If you are running background workloads at real-time prices today, switching to batch can reduce that cost by up to 50%

Example: 50,000 requests using Claude Opus 4.6 Assumes an average of 1,000 input and 500 output tokens per request.**

Metric	Rea-time Inference	Batch Inference
Input Cost (50M tokens @ \$5/M)	\$250.00	\$125.00
Output Cost (25M tokens @ \$25/M)	\$625.00	\$312.50
Total Cost	\$875.00	\$437.50

Pricing information current as of May 2026

By switching to Batch in this example, you save **\$437.50** on a single run. This enables you to use top-tier intelligence for massive data processing tasks that might otherwise be cost-prohibitive, while also creating new opportunities to optimize inference budgets across high-volume workloads.

Bypass Rate Limits

Batch jobs run on a dedicated throughput lane, completely separate from your real-time inference quota. Your production endpoints remain healthy while a 40,000-request batch job processes in the background across either provider. This helps reduce `429 Too Many Requests` errors in your data pipelines.

Asynchronous Processing

Submit a job and continue with other work. DigitalOcean manages the queue, retries, and delivery. You can poll for results when the job completes, or configure a webhook to receive notifications automatically (webhook delivery is coming soon).

Deeply Integrated with DigitalOcean

Batch inferencing is built into the DigitalOcean platform. Every part of the workflow, from file storage to job monitoring to usage analytics, runs on infrastructure you already use.

Powered by DigitalOcean Spaces

Input files (up to 200 MB) are uploaded directly to **DigitalOcean Spaces** via presigned URLs. There is no external storage to configure, no S3 buckets to provision, and no cross-account IAM policies to manage. The API generates a presigned upload URL, you `PUT` your `.jsonl` file, and Spaces handles the rest.

Results are delivered the same way. When a job completes, the results endpoint returns a presigned Spaces download URL. Result files are retained up to 30 days, so you can retrieve them on your own schedule.

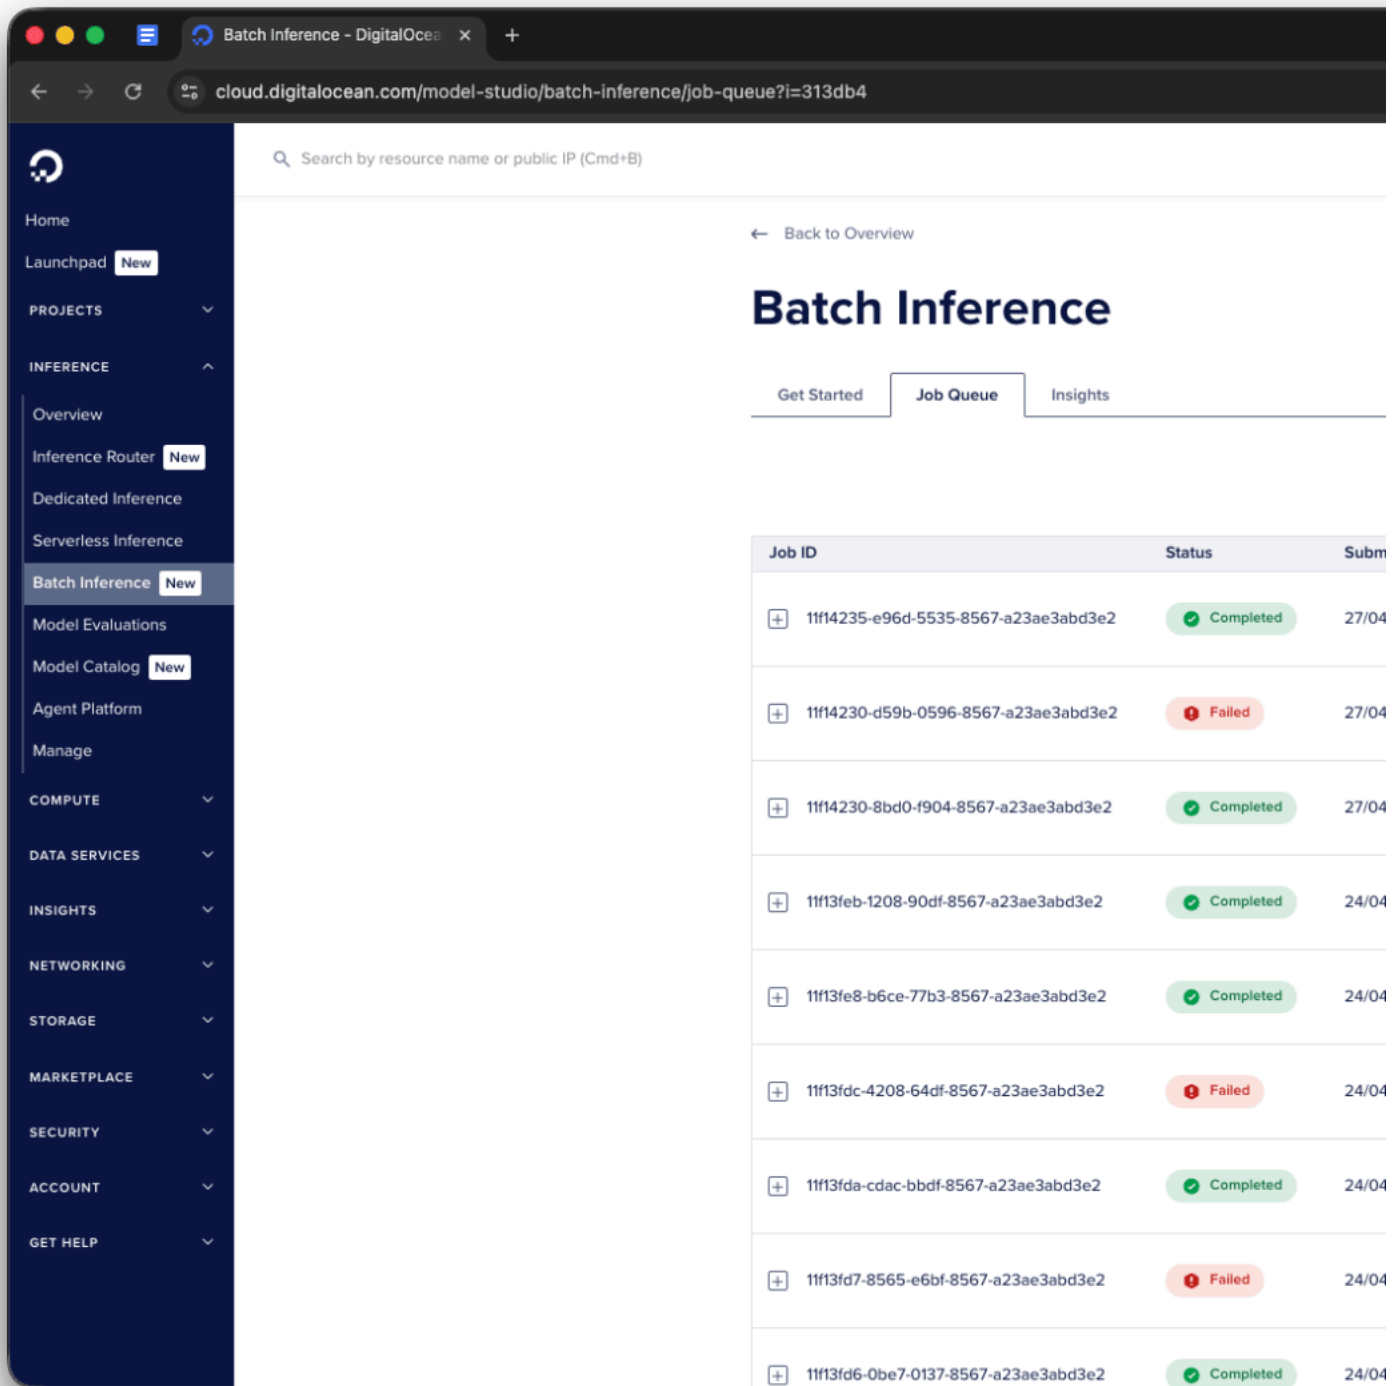
This is the same Spaces object storage that powers the rest of the DigitalOcean ecosystem, now integrated into your AI batch pipeline.

Job Queue: Track Every Job in Real Time

The [Batch Inference Job Queue](#) in the DigitalOcean Control Panel provides a live view of every batch job, with OpenAI and Anthropic jobs displayed side by side in a single list. For each job, you can view:

- **Status:** `awaiting_processing`, `in_progress`, `completed`, `failed`, `cancelled`
- **Progress:** total requests, completed, and failed counts, updated as the job runs
- **Timestamps:** when the job was submitted, started, and completed
- **Provider:** which provider is executing the batch

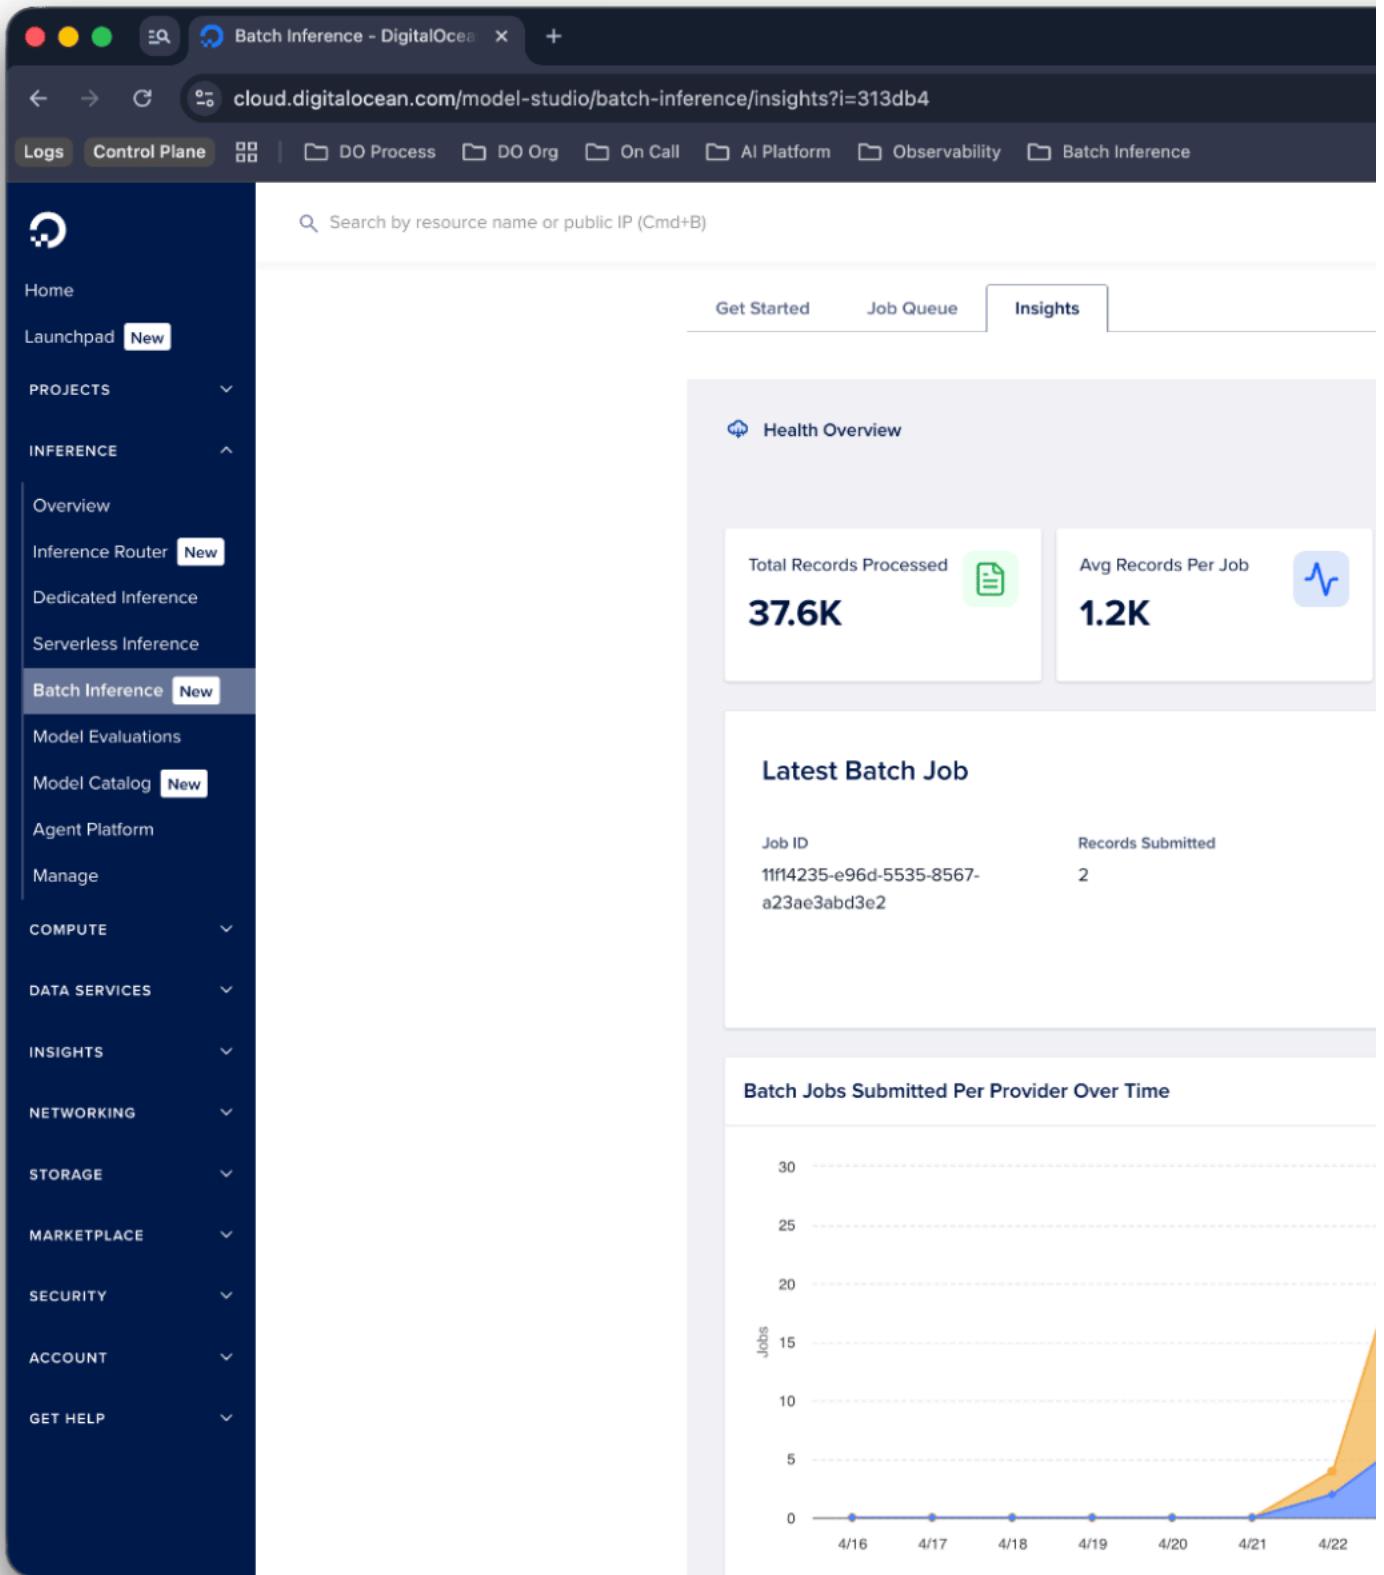
This eliminates the need to poll the API during development. You can monitor your jobs directly from the same Control Panel you use for Droplets, Databases, and Kubernetes.



Insights: Understand Your Usage

The [Batch Inference Insights](#) page provides a centralized view of batch usage across both providers. You can track token consumption, job volumes, and completion trends over time, all in one place rather than across separate OpenAI and Anthropic dashboards.

Use Batch Inference Insights to understand cost patterns, identify peak usage periods, and plan capacity for your batch pipelines.



Unified Billing

Token usage and job costs for both OpenAI and Anthropic batch workloads appear on a single DigitalOcean invoice. There are no separate bills to reconcile across providers and no additional payment methods to manage.

MCP Server Support

Batch Inference is also available as an [MCP \(Model Context Protocol\)](#) server, enabling seamless integration with AI-powered IDEs, agent frameworks, and any MCP-compatible client. This allows developers to create batch jobs, monitor their status, and retrieve results directly within their existing workflows.

Agents can be instructed to operate on input files, such as JSONL files for batch inference, by referencing a specified file path. Based on this context, the agent autonomously selects and invokes the appropriate MCP tools to handle file upload and initiate batch job creation. It can monitor status and upon completion, users can prompt the agent to retrieve the final job results and corresponding download URL, providing a seamless, end-to-end workflow with minimal manual intervention.

How It Works

The workflow is the same whether you are targeting OpenAI or Anthropic: **prepare, upload, submit, and retrieve**. All requests are sent to `https://inference.do-ai.run/v1` and authenticate with your Model Access Key.

- 1. Prepare your input file.** Create a `.jsonl` file where each line is a single inference request in the provider's native format. OpenAI lines include `custom_id`, `method`, `url`, and `body`. Anthropic lines include `custom_id` and `params`. The model is specified per request inside the file, giving you full flexibility within a single batch.
- 2. Upload your file.** Call `POST /v1/batches/files` with your file name to get a `file_id` and a presigned Spaces upload URL. Then `PUT` your `.jsonl` file to that URL. The presigned URL is valid for 15 minutes.
- 3. Create the batch job.** Call `POST /v1/batches` with your `file_id`, `provider` (`openai` or `anthropic`), and `completion_window` (24h). The endpoint, authentication, and response shape are identical for both providers. The only difference is the `provider` field.
- 4. Monitor and retrieve results.** Poll `GET /v1/batches/{batch_id}` for status, or monitor progress through the [Job Queue](#) in the Control Panel. Once the job reaches completed status, call `GET /v1/batches/{batch_id}/results` to get presigned download URLs for your output and error files. Result files are retained for **30 days**.

You can also **list all jobs** with `GET /v1/batches` and **cancel a running job** with `POST /v1/batches/{batch_id}/cancel`.

For full API details, code samples (cURL and Python), and input file format examples, see the [Batch Inference documentation](#).

Use Cases

Batch Inference is well-suited for any high-volume, non-latency-sensitive workload. The following examples are some of the most common patterns.

E-Commerce Catalog Enrichment

An e-commerce platform with 50,000 products needs SEO-friendly titles, marketing descriptions, and metadata tags for each listing. Rather than processing them through sequential API calls over several days, the entire catalog can be submitted as a single batch. You can use `gpt-4o-mini` for the English copy, then run a second batch through Claude for localized translations, all through the same pipeline with a different `provider` field.

Support Ticket Classification and Triage

Organizations can process a year's worth of support tickets in a single batch, classifying them by category, urgency, and sentiment while extracting structured fields like product name, issue type, and customer tier. The output is a clean `.jsonl` file ready to import into an analytics pipeline or CRM.

Content Moderation at Scale

Platforms with user-generated content, such as marketplaces, forums, and review sites, often need to scan thousands of posts, images, and listings for policy violations. Batch Inference allows you to process the entire backlog overnight without competing with your production moderation endpoint's rate limits.

Model Evaluation and Prompt Engineering

When developing a new system prompt, you can benchmark it against thousands of test cases by running the same evaluation suite against both OpenAI and Anthropic models through the same API. This enables side-by-side comparison of results at batch pricing, which is significantly lower than running the same evaluation in real time.

Document Processing and Data Extraction

Batch Inference can summarize thousands of legal contracts, research papers, or financial filings. It can also extract structured data such as dates, amounts, parties, and clauses from unstructured documents, or classify a backlog of invoices and receipts. These jobs can be large in volume but are rarely time-sensitive.

Getting Started

Batch Inference is available now on the DigitalOcean AI Platform.

Polling for job status is currently supported, with webhook notifications arriving soon for automated workflows. As the platform grows, expect expanded provider and model support.

The Bigger Picture

Inference has become the center of gravity in modern AI systems. Applications no longer make a single model call. They orchestrate multiple models, retrieve and synthesize data, execute tools, and repeat the cycle in production. This is a stack problem, not a model problem.

DigitalOcean's [AI-Native Cloud](#) was built to address exactly this. Five layers, one platform, open at every layer: GPU compute, inference, data and storage, agents, and the tools to connect them. Batch Inference is the latest addition to the inference layer, sitting alongside real-time Serverless Inference, the new Inference Router, Dedicated Inference, and a catalog of 25+ models across text, image, audio, and video.

Where real-time inference powers interactive experiences, Batch Inference handles the heavy lifting that happens behind the scenes. Together with GPU Droplets, Knowledge Bases, and Managed Databases (including Managed Weaviate (Private Preview) for vector workloads), they form a complete system for building production AI without stitching together services from multiple vendors.

The goal is to simplify the stack so you can focus on building.

Ready to get started? [Launch your first batch job](#) or visit the [product documentation](#) to learn more.

About the author(s)



snamdeo
Author
[See author profile](#)
[See author profile](#)



smirza
Author
[See author profile](#)
[See author profile](#)


Share



-
- [Product Updates](#)

Start building today

From GPU-powered inference and Kubernetes to managed databases and storage, get everything you need to build, scale, and deploy intelligent applications. [Sign up](#)



Related Articles




Product updates

OpenCode Now Supports DigitalOcean Inference Router for Intelligent Model Routing

[Musa Malik](#)

- May 28, 2026
- 3 min read

[Read more](#)






Product updates

Request-Based Autoscaling Is Now Generally Available on App Platform

Bikram Gupta

- May 22, 2026
- 4 min read

[Read more](#)




Product updates

Powering the Inference Era: Inside the DigitalOcean AI-Native Cloud

Vinay Kumar, Chief Product & Technology Officer

- May 4, 2026
- 7 min read

[Read more](#)



Company

- [About](#)
- [Leadership](#)
- [Blog](#)
- [Careers](#)
- [Customers](#)
- [Partners](#)
- [Referral Program](#)
- [Affiliate Program](#)
- [Press](#)
- [Legal](#)
- [Privacy Policy](#)
- [Security](#)
- [Investor Relations](#)

Products

- [GPU Droplets](#)
- [Bare Metal GPUs](#)
- [Inference Engine](#)
- [Data & Learning](#)
- [Droplets](#)
- [Kubernetes](#)
- [Functions](#)
- [App Platform](#)
- [Load Balancers](#)
- [Managed Databases](#)

- [Spaces](#)
- [Block Storage](#)
- [Network File Storage](#)
- [API](#)
- [Uptime](#)
- [Cloud Security Posture Management \(CSPM\)](#)
- [Identity and Access Management \(IAM\)](#)
- [Cloudways](#)
- [View all Products](#)

Resources

- [Community Tutorials](#)
- [Community Q&A](#)
- [CSS-Tricks](#)
- [Write for DOnations](#)
- [Currents Research](#)
- [DigitalOcean Startups](#)
- [Wavemakers Program](#)
- [Compass Council](#)
- [Open Source](#)
- [Newsletter Signup](#)
- [Marketplace](#)
- [Pricing](#)
- [Pricing Calculator](#)
- [Documentation](#)
- [Release Notes](#)
- [Code of Conduct](#)
- [Shop Swag](#)

Solutions

- [AI GPU Hosting](#)
- [H100 Cloud GPU](#)
- [AI Training GPU](#)
- [GPU Inference](#)
- [VPS Hosting](#)
- [Website Hosting](#)
- [VPN](#)
- [Docker Hosting](#)
- [Node.js Hosting](#)
- [Web Mobile Apps](#)
- [WordPress Hosting](#)
- [Virtual Machines](#)
- [View all Solutions](#)

Contact

- [Support](#)
- [Sales](#)
- [Report Abuse](#)
- [System Status](#)
- [Share your ideas](#)



Company



Products



Resources



Solutions




Contact



 © 2026 DigitalOcean, LLC. [Sitemap](#).

- 
- 
- 
- 
- 
- 
- 
- 
- 

 Dark mode is coming soon.

This site uses cookies and related technologies, as described in our [privacy policy](#), for purposes that may include site operation, analytics, enhanced user experience, or advertising. You may choose to consent to our use of these technologies, or manage your own preferences. Please visit our [cookie policy](#) for more information.

AGREE & PROCEED

DECLINE ALL

MANAGE CHOICES

[Cookie Preferences](#)