Published August 10, 2026 · Inaugural edition

America is naming its AI products long before they launch

Through the first seven months of 2026, Americans filed 408,346 U.S. trademark applications — 13.5% more than the same stretch of last year — and one in 21 of them claims artificial intelligence. Two-thirds of those AI applications are for products that aren't on the market yet, and the companies filing them are not the ones you would guess: LLCs and individuals account for nearly half, while Google ties for eleventh. Meanwhile the USPTO has cut the wait for a first Office Action from 251 days to 133. Here is the year so far.

By **GleanMark Research** August 2026 8 min read

JAN–JUL 2024	JAN–JUL 2025	JAN–JUL 2026
339,763	**359,823**	**408,346**
applications · 1.76% mention AI	applications · 2.87% mention AI	applications · 4.87% mention AI

The frame: seven months against seven months (January 1 – July 31), never against a full year. **All U.S. trademark applications** — those filed directly with the USPTO and those arriving through the Madrid Protocol, counted by the date each reached the Office rather than the date it was filed abroad, so all three years are complete and comparable. Two sections are narrower and say so where they appear: intent-to-use, because Madrid applications file under a different section of the Act that has no intent-to-use basis, and the refusal figures, because our office-action coverage on the Madrid route is still being rebuilt. Coverage of office actions on directly-filed applications in this window is 97.6%. Details at the end.

CONTENTS

01 DEMAND

The AI land grab — and who's really behind it

Start with the raw demand. Filings in the first seven months rose from 339,763 in 2024 to 408,346 this year — up 20% in two years — and artificial intelligence is doing far more than its share of the lifting. Applications whose goods-and-services text says "artificial intelligence" went from 5,965 to **19,867**, a 3.3× rise in two years, and from 1.76% of everything filed to **4.87%** — one in every 21 applications. Crypto, the last vocabulary to sweep the register, never got past 0.68% and is fading.

The more interesting question is *who* is filing, and the register answers it two ways. First, what they are filing on: **65.4% of AI applications are intent-to-use** — the basis you claim when the product isn't in commerce yet — against **48.2%** for every other application filed this year. That is a 17.2-point gap on a base of 18,761 AI applications filed directly with the USPTO: two out of every three AI names being claimed this year are attached to something that is not yet for sale. (This one measure covers directly-filed applications only. Madrid Protocol applications enter under a section of the Act that has no intent-to-use basis at all, so counting them here would add 1,106 filings that cannot be intent-to-use by definition.)

65.4% intent-to-use

> Share of 2026 AI filings made before the product is in commerce — versus 48.2% for everything else on the register.

Second, who is behind them — and it is not Big Tech. Break the AI filings down by the kind of applicant and **LLCs (30.5%) and individual people (15.3%) together make up almost half of them**; corporations of every size account for 36.6%. The league table of filers says the same thing more bluntly. The most active AI filer this year is OpenEvidence, with 62 applications. Tied for second, with 33 apiece, are Lemon Inc. and **one individual person in Oregon**, whose 33 applications were all filed intent-to-use inside five weeks. **Google ties for eleventh, with 19 — ten filers are ahead of it.**

The most active AI trademark filers, January–July 2026

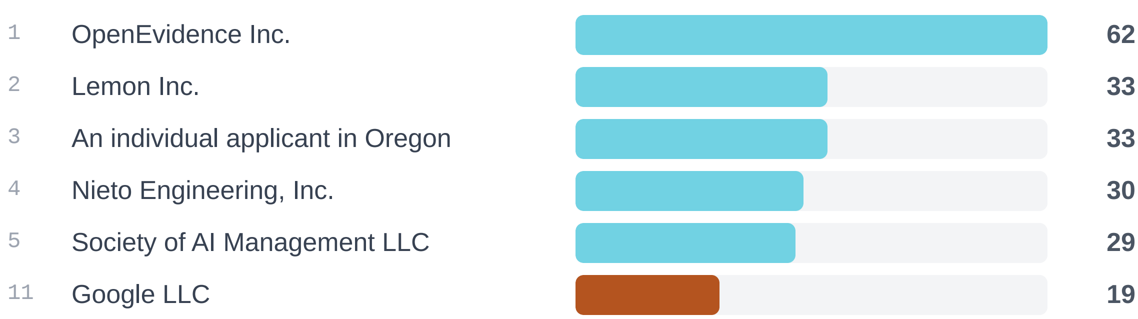

Applications whose goods/services text contains "artificial intelligence", by canonical owner name, directly-filed filings January 1 – July 31, 2026. The individual applicant is unnamed by choice — the filings are public record, but this is a data report, not a profile. Google ties for 11th with Cotiviti, which also filed 19. Source: GleanMark analysis of USPTO records.

Put the two cuts together and the shape of the AI land grab is clear: it is being run by small companies and individuals claiming names for products still in development — which is why the register reads as a pipeline of what is coming rather than a census of what has shipped. And they are mostly new to the register: only 35.6% of this year's AI applications came from an owner who had filed before — a shade under the 38.3% for everything else.

02 LANGUAGE

The word everyone uses and no one owns

Where does all that AI language land? Overwhelmingly in Class 042 — the software-and-technology-services class, itself up from 13,664 applications in the same seven months of 2021 to 45,142 today — and it has rewritten the class's vocabulary. Five years ago the words that filled Class 042 descriptions were the timeless furniture of tech filings: *management, consulting, research, analysis,* managing. "Intelligence" ranked 34th and "artificial" 46th. This year they rank **second and third**, behind only "management" — and behind it by 86 applications. **15,791 of the class's 45,142 applications now use the word "intelligence"**: better than one in three, against one in thirteen five years ago.

The most common words in software-class filings, then and now

JAN–JUL 2021		JAN–JUL 2026	
1	management	1	management
2	consulting	**2**	**intelligence**
3	research	**3**	**artificial**
4	analysis	4	analysis
5	managing	5	managing
6	mobile	6	consulting
7	monitoring	7	mobile
8	processing	8	research

Substantive words ranked by the NUMBER OF APPLICATIONS using them (not word occurrences, which double-count multi-class filings), in Class 042 goods/services descriptions, directly-filed applications filed January 1 – July 31 of 2021 vs 2026. The same stopword list —

application-form boilerplate such as "namely", "field", "temporary" and the generic filing vocabulary present in nearly every description — is applied to both periods, which is what makes the comparison fair. Source: GleanMark analysis of USPTO records.

And here is the punchline the register delivers on schedule: once every brand wants a word, the Office stops letting anyone own it. Among applications filed this year, **“AI” is the second-most-disclaimed term at the USPTO** — disclaimed 695 times by applicants formally renouncing any exclusive right to it. The only term ahead of it is “CO” — 697 — the abbreviation companies put at the end of their own names, as is “LLC” just below. Among words that are actually words, AI is now first, ahead of HEALTH (520), GROUP (461) and COFFEE (405). Five years ago “AI” ranked 169th, with 117 disclaimers — and those 2021 applications have had five more years to collect them. A disclaimer is the examiner’s way of saying a word belongs to the language, not to you. AI has become that kind of word.

> **169th → 2nd**
>
> Where “AI” ranks among the terms the USPTO makes applicants disclaim, comparing applications filed in the first seven months of 2021 and of 2026. Everyone may use it; no one may own it.

03 SUPPLY

Meanwhile, the USPTO is digging out — every clock is running faster

All of this demand lands on an agency that spent the early 2020s buried. The pandemic-era filing surge left the Office with a queue it has been working through ever since — and this year is the clearest evidence yet that it is winning. Every stage we can measure, measured by when the event actually happened, is faster than a year ago and much faster than two:

Three clocks, all running down

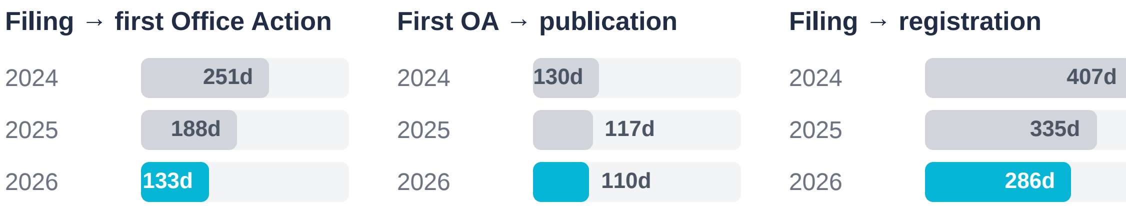

Median days, by the seven-month window the event occurred in, from the complete USPTO prosecution-event record (207K–213K first Office Actions per window; 102K–107K publications following an Office Action per window; registrations issued per window as shown). Filing→first Office Action is Office queue time; first OA→publication also includes the applicant's own response time. directly-filed applications only. Source: GleanMark analysis of USPTO records.

The headline is the first clock: an applicant whose first Office Action arrived in the first seven months of 2024 had waited a median **251 days** for it; for actions issued this year the median is **133 days**. Marks that hit an Office Action and fixed it reached publication 20 days sooner than in 2024. End to end, marks registering this year got there in a median **286 days, down from 407**, and the Office issued **250,888 registrations** in seven months — 47% more than in the same stretch of 2025. Faster medians *and* higher volume at the same time is what a backlog being cleared looks like, not a quiet year.

For applicants the practical read is simple: the years of “file and wait a year to hear anything” are ending. A mark filed today gets in front of an examiner in roughly four months — which also means conflicts, refusals and deadlines all arrive sooner.

04 THE GAUNTLET

What actually stops an application

Speed cuts both ways: more examination means more Office Actions, sooner. We hold the text of 231,117 of the 236,827 non-final actions the USPTO issued on directly-filed applications between January and July — **97.6%** — so what follows is close to the whole population rather than a sample of it. Across those and the final actions issued in the same window, **257,347 office actions carry at least one refusal or requirement**, spread over 233,640 applications. Four issues dominate, and they arrive in a near-tie:

The most common examiner objections, January–July 2026

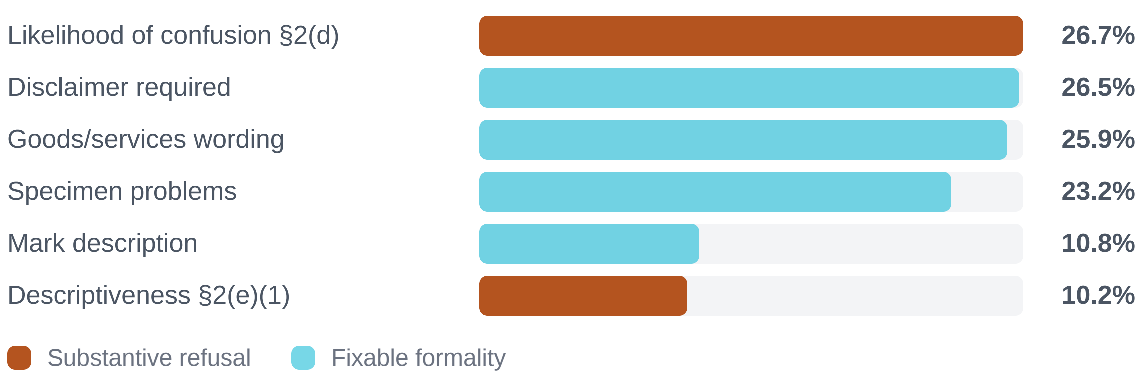

Share of the 257,347 office actions issued on directly-filed applications January 1 – July 31, 2026 carrying at least one refusal or requirement (grounds co-occur, so shares sum past 100%). Grounds are parser-derived; each ground shown here was checked against the examiner's own section headings. Advisory-only passages and the examiner's "no conflicting marks found" clause are excluded. Source: GleanMark analysis of USPTO office actions.

The pattern matters more than any single number. Only two of these — likelihood of confusion with an existing mark (§2(d), 26.7%) and descriptiveness (§2(e)(1), 10.2%) — are substantive judgments about whether the mark can be protected at all. The rest are **fixable formalities**: a disclaimer to add, goods-and-services wording to tighten, a specimen to re-file, a description to amend. Things a careful application avoids and a competent response cures. Most of the gauntlet is paperwork, not verdicts. The refusal to fear is §2(d) — the one you cannot paperwork your way out of, and the one the next section is about.

05 THE BLOCKERS

The marks that block the most Americans are words you'd think nobody could own

When an examiner refuses a mark under §2(d), they name the registration standing in the way. Tally those citations by the number of different applicants each registration stopped — not by applications, which one busy filer can inflate on its own — and the top of the list holds no tech giants and no famous brands. It is a spirit, a sport, a profession and a country's coffee.

#1 ***TEQUILA*** Consejo Regulador del Tequila, A.C. · registered 2017 **49 applicants blocked**

The certification mark for the spirit itself, held by the Mexican council that certifies tequila. It took fourteen years to register — filed 2003, granted 2017 — and this year the agave boom sent 49 different applicants into it.

#2 ***PICKLE-BALL*** Pickleball Holdings LLC · registered 1974 **46 applicants blocked**

A 1974 registration on the sport's own name, covering paddles, nets and balls sold as a unit. Pickleball brands keep arriving at the register, and 46 of them walked into the same fifty-year-old wall this year.

#3 ***REALTORS*** National Association of Realtors · registered 1949 **17 applicants blocked**

The profession’s own name, registered in 1949 and policed ever since. Its siblings REALTOR (1950) and REALTOR.COM (2012) blocked another 16 and 11 applicants — three of this year’s five most obstructive registrations belong to one trade association.

#4 ***AGENTIC*** Epic Games, Inc. · registered 2023 **9 applicants blocked**

The AI industry’s new favorite adjective has belonged to the maker of Fortnite since 2023, for software development tools. Two applications used the word “agentic” in all of 2023; 513 have been filed in the first seven months of this year — and nine of those applicants have already been refused over Epic’s two registrations.

#5 ***COLOMBIAN*** The Republic of Colombia · registered 1981 **8 applicants blocked**

A certification mark on the word “Colombian” — for coffee, owned by the country itself since 1981.

One more registration, because it shows why the ranking method matters. The mark cited against the most applications this year is the two-letter AI, registered in December 2024 by the State University of New York for computer-science courses — 77 citations, half again as many as TEQUILA. But all 77 trace to just two applicants, each filing a whole family of AI-something names, and by applicants blocked it does not come near this list. A registration that stops two determined filers seventy-seven times is a different animal from one that stops forty-nine different businesses once each. Only the second is a wall the market keeps walking into.

THE BOTTOM LINE

Seven months of 2026 read like an economy naming its next act before the curtain rises: a wave of small companies and individuals staking AI names for products that mostly don't exist yet, a word so common the Office now makes applicants renounce it, and an agency that has halved its examination wait while filings climbed 13.5%. **Trademark filings are intent, filed in public — and right now America's intent is AI, filed by newcomers, and arriving at an Office clearing the queue faster than at any point in the three years we can measure.**

METHODOLOGY & DATA

How we measured it

Figures come from GleanMark's structured copy of the public USPTO record, queried directly. Each instrument is labeled where it is used, and each limit is stated rather than buried:

Scope — all U.S. trademark applications

This report counts BOTH applications filed directly with the USPTO and applications arriving through the Madrid Protocol (serial numbers beginning 79). Madrid applications are counted by the date they reached the USPTO — their first entry in the Office's own transaction record — NOT by the date they were filed abroad. That distinction decides whether the comparison is honest: Madrid applications land on the U.S. register months after they are filed internationally, so counting them by filing date shows 1,798 arriving in January 2026 and 49 in July, a collection curve rather than a real decline. Counted by arrival, the same seven months hold steady — 13,898 in 2024, 15,282 in 2025, 14,877 in 2026 — and all three years are complete. TWO sections stay narrower, each for a stated reason. INTENT-TO-USE covers directly-filed applications only, because Madrid applications enter under §66(a), which has no intent-to-use basis at all; their flag is false by construction, so including them would dilute a real finding with a definitional zero. THE REFUSAL AND CITATION FIGURES likewise cover directly-filed applications only, because our coverage of office actions issued on the Madrid route is currently under 1%, against 97.6% for directly-filed ones. The Class 042 vocabulary comparison is also directly-filed, so its 2021 and 2026 baselines are built the same way.

Filing demand

Every directly-filed U.S. application by filing date, January 1 – July 31 of 2024, 2025 and 2026, compared like-for-like — seven months against seven months, never against a full year. AI and crypto terms are case-insensitive matches in the goods/services text: for AI, the literal phrase "artificial intelligence" only. Bare "AI" is deliberately excluded, so the AI share here is a floor, not a ceiling. Intent-to-use is the application's own ITU basis flag; "repeat filer" means the exact canonical owner name appears on an application filed before this year, which undercounts repeats for both cohorts equally (name variants), so the AI-vs-rest comparison holds even though both levels are floors.

USPTO speed

Medians from the complete prosecution-event history: "first Office Action" is the earliest non-final action mailed or e-mailed per application; "publication" is the published-for-opposition event. Each metric is grouped by the seven-month window the event happened in, so a period reflects the Office's speed at that time rather than the fate of a filing cohort. Filing→first OA measures Office queue time; first OA→publication includes the applicant's own response time; filing→registration covers marks whose registration issued in the window. Cohort sizes: 207K–213K first Office Actions per window; 102K–107K publications following an Office Action per window.

Refusal mix & blocking registrations

GleanMark's parsed office-action corpus: 257,347 directly-filed office actions dated in this window that carry at least one refusal or requirement — non-final actions (the great majority) and final actions. We hold the text of 97.6% of the 236,827 non-final actions the Office issued in the window (231,117), so these are near-population figures rather than a sample — but they are still parser-derived, and every ground reported above was checked against the examiner's own section headings before publication. Advisory-only passages and the examiner's "search of office records — no conflicting marks found" clause are excluded; counting those as refusals inverts their meaning. Blocking registrations are ranked by the number of distinct applicants refused over them under §2(d), not by applications: one applicant filing a family of forty names produces forty citations and one blocked business. Counts were measured on 3 August 2026; the corpus grows as documents are parsed, so a later re-run returns marginally higher absolute numbers at the same shares.

Why there is no year-over-year refusal trend here

Because we cannot honestly show one yet. Our coverage of domestic office actions was 97.6% in this window but 61.3% in the same seven months of 2025 — so any "increase" in refusals or citations between those two years would mostly be our own pipeline catching up, not examiners changing behavior. Every refusal and blocking-registration figure on this page is 2026 only. The filing-side figures come straight from the register and are comparable across years, which is why those are the only trends shown.

Vocabulary & disclaimers

Class 042 word ranks compare goods/services text of applications FILED in January 1 – July 31 2021 vs 2026, with one shared stopword list (form boilerplate plus the generic vocabulary present in nearly every description) applied to both periods. The disclaimer comparison is by RANK among disclaimed terms of applications filed in each period — ranks are comparable across periods even though older cohorts have accumulated more disclaimer entries.

What we did not measure

How registered marks overcame each refusal type needs response-side outcome data our corpus does not yet support cleanly; it is on the roadmap. A filing is intent, not a granted right, and an office action is an examiner's opening position, not a verdict. Chart data, term lists and the serial numbers behind every figure are available to reporters on request.

ABOUT GLEANMARK

The research platform behind this report

GleanMark is a U.S. trademark intelligence platform for attorneys and law firms, built on a structured copy of the public USPTO record — about 14 million applications and registrations, 648,000+ TTAB proceedings, and 240 million+ prosecution events, refreshed continuously. Every trademark, owner, law firm, correspondent and TTAB proceeding on the register has a research profile, and the same corpus powers GleanMark's monthly and quarterly data reports, its evidence-based clearance reports, and its office action response drafter. GleanMark is live with paying customers; pricing is published at gleanmark.com/pricing.

CITE THIS REPORT

"GleanMark, *State of Trademark Filings — January–July 2026*."

https://gleanmark.com/reports/state-of-trademark-filings-jan-jul-2026

Download the PDF · Press & data requests, including full methodology and chart data: contact us.

This is the **inaugural edition** of a recurring study. **This page is permanent** — the figures below will not be revised or replaced, so a citation to this URL resolves to these numbers indefinitely. The next edition covers the second half of 2026 and publishes as a separate report in January 2027.

READ THE QUARTERLY DEEP DIVES

PROSECUTION

Final Refusals Fell Every Month of Q2 Even as Examiners Issued More Non-Finals

The examining corps split in two directions this spring — inside the refusal data.

TTAB

Zara, Nike, and a Wave of Big-Tech Watch-Filings: TTAB Disputes Climbed 5.5% in Q2

Who is fighting hardest at the Trademark Trial and Appeal Board — and how those disputes end.

FILING TRENDS

Meta's New Smart Glasses Show Up on the Trademark Register the Same Day It Launched Them

June's filing trends, from product launches to the month's most active applicants.