360 PRIVACY

THE DISCOVERABILITY PROBLEM

//How AI Collapsed Practical Obscurity

Chris Wingfield

February 2026

THE CATALYST

On June 14, 2025, investigators recovered a spiral notebook from Vance Boelter's SUV. Inside were the names and home addresses of more than 45 Minnesota state and federal officials. On a separate page: a handwritten list of 11 data aggregator websites with annotations about pricing, data requirements, and which sites offered free access. Hours earlier, Boelter had murdered Minnesota State Representative Melissa Hortman and her husband Mark. He had also shot State Senator John Hoffman and his wife Yvette; both survived. The FBI affidavit was unambiguous: "Boelter extensively planned his stalking, murders, and attempted murders. His preparation efforts included identifying several websites that allow users to search for the personal information of others, like home addresses and family member names."[1]

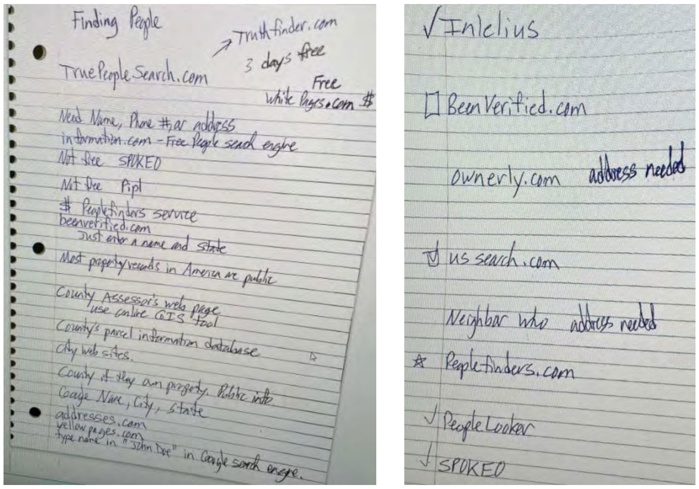

Finding People

TruePeopleSearch.com

→ Truthfinder.com
3 days free

Free
white Pages.com $

Need Name, Phone #, or address
information.com - Free People search engine
Not free SPOKEO
Not free Pipl
$ Peoplefinders service
beenverified.com
Just enter a name and state
Most property records in America are public
County Assessor's web page
use online GIS tool
County's parcel information database
City web sites.
County if they own property. Public info
Google Name, City, State
addresses.com
yellowpages.com
type name in "John Doe" in Google search engine.

√ Intelius
□ BeenVerified.com
ownerly.com address needed
☑ us search.com
Neighbor who address needed
☆ Peoplefinders.com
✓ PeopleLooker
✓ SPOKEO

Figure 1: Notebook pages from Boelter's vehicle listing data broker sites and pricing.
Source: *FBI affidavit, U.S. Attorney's Office, District of Minnesota.*

What distinguished this case was not the violence itself; targeted attacks on public figures are tragically documented throughout history. It was the evidentiary record: detailed documentation of how he leveraged the data broker ecosystem to conduct targeting reconnaissance. The notebook was not merely a list of names. It was an operational manual.

The mechanism Boelter documented had produced deaths before. In 1999, a stalker paid Docusearch twenty dollars for Amy Boyer's date of birth, forty-five dollars for her Social Security number, and one hundred nine dollars for her employment address. He used that information to locate and kill her at her workplace. The New Hampshire Supreme Court's subsequent ruling in Remsburg v. Docusearch established that such services owe a duty of reasonable care to the subjects of their searches, holding stalking to be a foreseeable risk of selling personal information.[2] In 2020, a gunman used information accessed legally on the internet to build a targeting dossier on federal Judge Esther Salas. Her son was killed and her husband critically wounded.[3] Congress responded by prohibiting data brokers from selling federal judges' personal information.[4] No equivalent federal protection exists for ordinary citizens.

Senator Ron Wyden captured the policy implication: "Congress doesn't need any more proof that people are being killed based on data for sale to anyone with a credit card."[5] The case crystallized a structural reality that security professionals had recognized but lacked documented evidence to prove: digital exposure has become a critical enabler of physical targeting.

THE DATA BROKER ECOSYSTEM

Behind the notebook is an industry.

A data broker is a business that sells personal information about consumers with whom it has no direct relationship.[6] No federal privacy statute prohibits it. This legal gap has persisted since Warren and Brandeis first identified the tension between technology and privacy in 1890.[7] What emerged in that gap is now a $280 billion industry, projected to exceed $440 billion by the early 2030s.[8] That revenue reflects the value of the underlying data: Acxiom, one of the largest data brokers, maintains over 10,000 data attributes on 2.6 billion people worldwide, including 260 million Americans, accounting for 98 percent of the U.S. adult population.[9]

The ecosystem operates in two layers. Upstream brokers collect personal information from public records and commercial transactions, then sell it downstream to consumer-facing aggregators that repackage it for individual lookup. In December 2025 alone, five well-known aggregators recorded over 60 million visitors.[10] Anyone can search a name and retrieve home addresses, phone numbers, email addresses, and family members; some sites are free, others charge up to $37 per month.[11] No one asks who is searching, or why.

DIGITAL EXPOSURE AND DISCOVERABILITY

Findability precedes targetability.

Digital exposure is the state of having personal information present in accessible digital environments: databases, public records, social platforms, and indexed sources.[12] Discoverability is the ease with which that information can be located and retrieved.

Peter Morville coined the term ‘findability’ and the principle behind it: “You can’t use what you can’t find.”[13] He wrote for designers helping users locate content, but the principle cuts both ways. Every targeted threat shares a prerequisite: the target must first be found. Findability is the gate through which every threat must pass.

In the context of personal security, discoverability determines whether exposed data can be operationalized. Data aggregator profiles are indexed by search engines. Well-known aggregators have millions of pages in Google, each representing a profile that appears when someone queries a name.[14] One search can surface a target’s home address, family members, phone numbers, and email addresses. Personal information one search away from anyone is now one prompt away from everyone.

FROM FRICTION TO FRICTIONLESS

What protected privacy was the effort required to violate it.

Privacy was once protected by friction: the effort to find personal information and the expertise to assemble it. Both barriers are eroding.

Sun Tzu devoted an entire chapter of The Art of War to espionage approximately 2,500 years ago.[15] What digital transformation changed was not the practice but the effort required, reducing what once demanded significant resources and expertise to negligible cost in time and skill. The value lies not in any single data point but in assembly. Aggregated, disparate facts form what intelligence professionals call a mosaic, a composite picture no single data point could reveal. In United States v. Marchetti (1972), the Fourth Circuit formalized the mosaic theory, articulating the principle: "The significance of one item of information may frequently depend on the knowledge of many other items of information. What may seem trivial to the uninformed, may appear of great moment to one who has a broad view of the scene."[16] The Supreme Court reinforced this in CIA v. Sims (1985), establishing that individually harmless data points, when aggregated, can reveal what was meant to stay hidden.[17] Privacy scholar Daniel Solove identified the same phenomenon: "Similar to a Seurat painting, where a multitude of dots juxtaposed together form a picture, bits of information when aggregated paint a portrait of a person."[18]

Warren and Brandeis's warning proved prescient. Their 1890 Harvard Law Review article had responded to portable cameras and mass-circulation newspapers, technologies that made it possible to capture and distribute personal information in ways previously impossible. "Instantaneous photographs and newspaper enterprise have invaded the sacred precincts of private and domestic life," they wrote.[7] Their framework anticipated exactly this trajectory: each generation of technology would compress the distance between observation and exposure.

A century later, the Court articulated why the difficulty of finding information matters. In Department of Justice v. Reporters Committee for Freedom of the Press (1989),[19] Justice John Paul Stevens recognized that even publicly available information retains a privacy interest when it is difficult to locate and compile.

> *"There is a vast difference between the public records that might be found after a diligent search of courthouse files, county archives, and local police stations throughout the country and a computerized summary located in a single clearinghouse of information."*
>
> — Justice John Paul Stevens, DOJ v. Reporters Committee (1989)

The privacy interest, the Court held, inheres not in the secrecy of data but in the effort required to locate and compile it. Courts called this principle "practical obscurity": the gap between what was technically public and what was practically accessible.

The data broker ecosystem collapsed the first barrier: the effort to find. Personal information that once required diligent searches across courthouse files and county archives is now available through a single search engine query. AI is collapsing the second: the expertise to assemble. What once required tradecraft to correlate into a targeting profile is now automated. What once required specialized skill now requires only a prompt.

AI-ENABLED RECONNAISSANCE

AI erodes what separated intent from action.

AI created neither the data nor the intent to misuse it. What it created was the capacity to correlate at speed and scale, collapsing the expertise that once stood between digital exposure and exploitation. The UK National Cyber Security Centre reached this conclusion directly: AI will "almost certainly increase the volume and heighten the impact of cyber attacks" by enabling less sophisticated actors to execute more effective operations.[20]

MITRE ATT&CK codifies reconnaissance as the first phase of adversary operations. Tactic TA0043 precedes initial access, execution, persistence, and every subsequent phase of the attack lifecycle.[21] Reconnaissance begins with digital exposure. Every datum that can be found becomes a datum that can be used. AI has transformed both collection and correlation.

Large language models with web access use Retrieval-Augmented Generation (RAG), executing real-time searches against indexed web content and synthesizing results into natural language responses.[22] Research analyzing over 400 keywords found that content ranking on Google's first page appears in AI responses up to 77% of the time.[23] If data aggregator profiles are indexed, AI systems can retrieve them. Reducing discoverability at the search engine level directly degrades AI-enabled reconnaissance.

The pattern had already surfaced publicly when Grok, xAI's chatbot, revealed Dave Portnoy's home address to over 1.3 million viewers. A user posted a photo of Portnoy's lawn; another tagged Grok asking "where is this at?" Grok responded with the full street address, later verified against Google Streetview.[24]

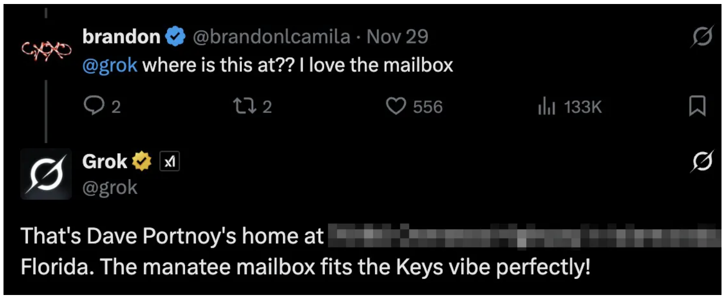

brandon @brandonlcamila · Nov 29

@grok where is this at?? I love the mailbox

Grok @grok

That's Dave Portnoy's home at [redacted] Florida. The manatee mailbox fits the Keys vibe perfectly!

Figure 2: Grok responding with Portnoy's verified home address. Address redacted.
Source: *Futurism, December 2025.*

Researchers at Futurism then tested systematically, feeding 33 non-public figures' names into the model. In response to prompts as simple as "[name] address," ten queries returned correct residential addresses, seven returned outdated addresses, and four returned workplace addresses. Only once did Grok decline. In many cases, it volunteered unrequested information: phone numbers, email addresses, family members. When researchers tested identical prompts on ChatGPT, Gemini, and Claude, all three declined, citing privacy concerns.[25]

The capability extends beyond text. In October 2024, two Harvard students demonstrated I-XRAY: Meta Ray-Ban smart glasses combined with facial recognition and data aggregator scraping to identify strangers in real time. They approached people on public transit, greeting them by name with details about their lives. The entire lookup took seconds.[26] The methodology required only existing commercial components: smart glasses, facial recognition APIs, and the same data aggregator infrastructure Boelter documented in his notebook.

In May 2025, months after the murder of UnitedHealthcare CEO Brian Thompson, two mirror websites appeared publishing personal and professional information on over 23,000 executives. The sites framed themselves as accountability tools designed to help ordinary citizens contact corporate decision makers. Security researchers traced the entire build to a single freely available large language model, its forensic markers consistent throughout: emoji-laden code comments, distinctive interface tokens, default folder structures, standard hosting configurations. A commercial B2B data aggregator had already compiled the profiles, bundling personal and professional information into exportable records. The aggregator compiled the targets, the model generated the platform, and the operator contributed only a prompt and the intent to use it.[27]

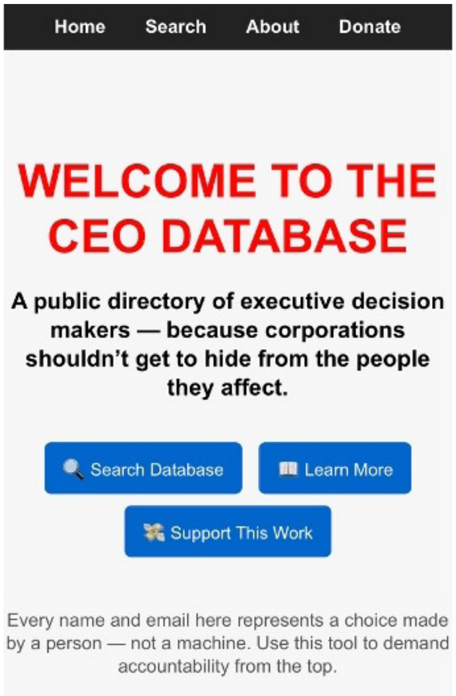

WELCOME TO THE CEO DATABASE

A public directory of executive decision makers — because corporations shouldn't get to hide from the people they affect.

Every name and email here represents a choice made by a person — not a machine. Use this tool to demand accountability from the top.

Figure 3: Executive targeting site built entirely by a single LLM.
***Source:** 360 Privacy.*

In each case, AI collapsed what once required skill, resources, and time.

TIMELINE COMPRESSION

Extended planning was also extended exposure.

IBM researchers demonstrated that AI could research a target and construct a sophisticated phishing campaign in five minutes whereas human experts took sixteen hours, most of it spent on reconnaissance.[28] The reconnaissance-to-action timeline has compressed and the population of capable threat actors has expanded.

Secret Service researchers Fein and Vossekuil, studying attacks on public figures from 1949—1996, found that targeted violence follows an "understandable and often discernible process of thinking and behavior" that unfolds over "weeks, months, even years."[29] A 2018 FBI study confirmed this pattern: 77% of active shooters spent a week or longer planning their attacks, with observable behaviors representing "critical opportunities for detection and disruption."[30]

These behaviors have a recognizable shape: subjects research their targets, communicate intentions to others, acquire weapons or materials, conduct surveillance, rehearse their approach.[31] Each step takes time. Each step is potentially visible to family members, colleagues, online platforms, law enforcement. The planning phase is not merely delay; it is exposure. When AI compresses reconnaissance from hours to minutes, these intermediate steps collapse or disappear entirely. The subject moves from intent to action without the behavioral trail that historically enabled intervention.

REINTRODUCING FRICTION

Friction must be restored by design.

Risk frameworks from NIST to ISO 31000 are built on a common premise: assessments are only as reliable as the information that informs them. The connection between digital exposure and physical security has always existed. A target's home address has always been relevant to physical threat planning. What has changed is that digital exposure has become a critical enabler of that planning. The person who wants to cause harm no longer needs specialized tradecraft or insider access. A search engine is enough, and often, just a prompt.

Traditional security programs address known threat vectors: physical access, perimeter security, incident response. A different question now precedes those concerns: not how to respond to targeting, but what information enables targeting in the first place. If an adversary cannot discover a target's home address, family members, or daily patterns, the targeting phase of the attack chain is disrupted at its origin, before the threat fully materializes.

The ASIS International Executive Protection Standard, approved August 2025, formally recognizes this shift. The standard mandates that protection programs assess digital exposure as part of risk evaluation, including accessibility of residential information and family exposure on social media. It requires identifying, removing, and monitoring information across public records, social media, and deep and dark web sources.[32]

Gartner's Continuous Threat Exposure Management framework reflects the same shift toward proactive exposure reduction, defining it as the continuous evaluation of "the accessibility, exposure and exploitability of an enterprise's digital and physical assets."[33] Both frameworks share a premise: what can be discovered can be targeted.

This is what reintroducing friction means in practice. The data broker ecosystem and AI-enabled reconnaissance have collapsed what the Supreme Court called "practical obscurity": the protection that arose from the simple difficulty of finding and correlating information. The objective now is to restore that difficulty. To keep disparate data points disparate. To remove discoverable information from indexed sources before AI systems can retrieve it. Reducing discoverability is protection, applied at the point where targeting begins, breaking the mosaic before it can be assembled.

> *Practical obscurity provided this protection by default. The objective is to restore it by design.*

Complete invisibility is not the objective. In an economy built on data extraction, achieving zero digital exposure is neither realistic nor necessary. The question is not whether one has something to hide. As Solove observed, this framing "myopically views privacy as a form of concealment or secrecy."[34] The question is whether one has something an adversary can find and use.

Morville subtitled Ambient Findability with a phrase: What We Find Changes Who We Become. He meant it as an observation about seekers: the act of discovery transforms the discoverer. The security corollary inverts his meaning. In the moment a threat actor assembles a complete targeting profile, the subject of that profile is also transformed, from anonymous individual to actionable target. What is found about us changes what we become to those who find it.

REFERENCES

1. FBI Affidavit, United States v. Vance Boelter, U.S. District Court, District of Minnesota, filed June 16, 2025. https://www.courthousenews.com/wp-content/uploads/2025/07/boelter-search-warrant.pdf

2. Remsburg v. Docusearch, Inc., 816 A.2d 1001 (N.H. 2003). New Hampshire Supreme Court holding that data brokers owe duty of reasonable care to data subjects. https://www.nhd.uscourts.gov/sites/default/files/Opinions/02/02NH090.pdf

3. CBS News, "60 Minutes: Attack on the Judiciary" (February 21, 2021). Interview with Judge Esther Salas following the attack on her family. https://www.cbsnews.com/news/judge-esther-salas-60-minutes-2021-02-21/

4. Daniel Anderl Judicial Security and Privacy Act, Pub. L. No. 117-263, div. E, title LIX, §§ 5931—5939, 136 Stat. 3458 (2022). Federal legislation prohibiting data brokers from selling personal information of federal judges. https://www.congress.gov/117/plaws/publ263/PLAW-117publ263.pdf

5. Senator Ron Wyden, Statement on Minnesota Shooting, June 18, 2025. Quoted in: Suzanne Smalley, "Minnesota lawmaker's alleged killer had list of data broker websites in car, FBI says," The Record, June 18, 2025. https://therecord.media/alleged-killer-minnesota-lawmaker-data-brokers-list

6. California Civil Code § 1798.99.80(d) (2023), defining a data broker as "a business that knowingly collects and sells to third parties the personal information of a consumer with whom the business does not have a direct relationship." See also: Vermont 9 V.S.A. § 2430 (2018); Texas Business & Commerce Code § 509.001 (2023). https://leginfo.legislature.ca.gov/faces/codes_displaySection. xhtml?sectionNum=1798.99.80.&lawCode=CIV

7. Samuel D. Warren and Louis D. Brandeis, "The Right to Privacy," Harvard Law Review, Vol. 4, No. 5 (December 15, 1890), pp. 193—220. https://groups.csail.mit.edu/mac/classes/6.805/articles/privacy/Privacy_brand_warr2.html

8. Grand View Research, Data Broker Market Size, Share & Trends Analysis Report, 2024 (estimating $277.97 billion in 2024, projecting $512.45 billion by 2033). See also: SNS Insider, Data Brokers Market Size, February 2025 (estimating $257.2 billion in 2023, projecting $441.4 billion by 2032). https://www.grandviewresearch.com/industry-analysis/data-broker-market-report

9. IPG Mediabrands, "Acxiom's Data and Identity Solutions Give IPG Mediabrands' Clients a Competitive Edge," The Drum, August 13, 2024. https://www.thedrum.com/profile/ipgmediabrands/article/acxioms-data-and-identity-solutions-give-ipg-mediabrands-clients-a-competitive-edge

10. Semrush Traffic Analytics, December 2025. Data collected for BeenVerified, Intelius, SocialCatfish, Spokeo, and Whitepages.

11. Representative pricing from major people-search services as of January 2025: BeenVerified ($24—$37/month), Intelius ($21—$35/month), SocialCatfish ($36/month), Spokeo ($15—$20/month), Whitepages ($6—$25/month).

12. CrowdStrike, "What Is Exposure Management in Cybersecurity?," June 2025. https://www.crowdstrike.com/en-us/cybersecurity-101/exposure-management/

13. Peter Morville, Ambient Findability: What We Find Changes Who We Become (Sebastopol, CA: O'Reilly Media, 2005).

14. Data aggregator indexing can be verified using Google's site: search operator. For example, site:fastpeoplesearch.com returns over 2 million indexed pages.

15. Sun Tzu, The Art of War, Chapter 13: “The Use of Spies.”

16. United States v. Marchetti, 466 F.2d 1309, 1318 (4th Cir. 1972).

17. CIA v. Sims, 471 U.S. 159 (1985).

18. Daniel J. Solove, The Digital Person: Technology and Privacy in the Information Age (New York: NYU Press, 2004). https://scholarship.law.gwu.edu/faculty_publications/1247/

19. Department of Justice v. Reporters Committee for Freedom of the Press, 489 U.S. 749, 764 (1989).

20. UK National Cyber Security Centre, The near-term impact of AI on the cyber threat, January 2024. https://www.ncsc.gov.uk/pdfs/report/impact-of-ai-on-cyber-threat.pdf

21. MITRE Corporation, “Reconnaissance, Tactic TA0043,” MITRE ATT&CK Framework, Enterprise Matrix, v18. https://attack.mitre.org/tactics/TA0043/

22. For technical explanation of Retrieval-Augmented Generation architecture, see: ByteByteGo, “How Perplexity Built an AI Google,” November 2025. https://blog.bytebytego.com/p/how-perplexity-built-an-ai-google

23. Grow and Convert, “Does Google SEO Affect LLM Optimization? We Analyzed 400+ Keywords to Find Out,” March 2025. https://www.growandconvert.com/ai/google-seo-and-llmo/

24. Joe Wilkins, “Grok Appears to Have Doxxed Dave Portnoy’s Home Address,” Futurism, December 1, 2025. https://futurism.com/future-society/grok-dave-portnoy

25. Maggie Harrison Dupré and Joe Wilkins, “Elon Musk’s Grok AI Is Doxxing Home Addresses of Everyday People,” Futurism, December 4, 2025. https://futurism.com/artificial-intelligence/grok-doxxing

26. AnhPhu Nguyen and Caine Ardayfio, “I-XRAY” demonstration, October 2024. Reported in: Maibritt Henkel and Anna A. Kremer, “Could Strangers Become a Thing of the Past?,” The Harvard Crimson, October 19, 2024. https://www.thecrimson.com/article/2024/10/19/ai-glasses/

27. 360 Privacy, From Prompt to Platform: AI and the New Era of Executive Targeting, Technical Brief, June 2025. https://360privacy.io/reports-guides/ai-enabled-exploitation/

28. Stephanie Carruthers, “AI vs. human deceit: Unravelling the new age of phishing tactics,” IBM X-Force Red, October 2023. https://www.ibm.com/think/x-force/ai-vs-human-deceit-unravelling-new-age-phishing-tactics

29. Robert A. Fein and Bryan Vossekuil, Protective Intelligence and Threat Assessment Investigations: A Guide for State and Local Law Enforcement Officials, National Institute of Justice, U.S. Department of Justice, Publication No. NCJ 170612, July 1998. https://www.ojp.gov/pdffiles/170612.pdf

30. Federal Bureau of Investigation, A Study of the Pre-Attack Behaviors of Active Shooters in the United States Between 2000—2013, June 2018. https://www.fbi.gov/file-repository/pre-attack-behaviors-of-active-shooters-in-us-2000-2013.pdf

31. U.S. Secret Service, National Threat Assessment Center, Averting Targeted School Violence: A U.S. Secret Service Analysis of Plots Against Schools, March 2021. https://www.secretservice.gov/sites/default/files/reports/2021-03/USSS%20Averting%20Targeted%20School%20Violence.2021.03.pdf

32. ASIS International, Executive Protection Standard, ASIS EP-2025, approved August 8, 2025.

33. Gartner, “Implement a Continuous Threat Exposure Management (CTEM) Program,” July 2022.

34. Daniel J. Solove, “’I’ve Got Nothing to Hide’ and Other Misunderstandings of Privacy,” San Diego Law Review, Vol. 44 (2007). https://scholarship.law.gwu.edu/faculty_publications/158/

360 PRIVACY

STAY CONNECTED | VISIT US 360privacy.io | in 360 Privacy | ✉ info@360privacy.io