

Manufactured Trust

Inside SEO poisoning, hacked websites, fake Google reviews and political bot farms.



Original illustration · An established identity can carry a message its owner never approved.

Who controls the machines and accounts? Who pays for the results? How do ordinary searches become entry points... and what happens when investigators catch the people behind them?

This source-linked black hat SEO case study follows recent exposures and court cases, separating observed activity from sales claims, allegations and measurable impact.

Focus: organic manipulation. Advertising is excluded.

Updated 13 October 2026 Historical context + 2026 disclosures Evidence-linked reference edition

1,800+

Compromised Windows web servers reported in Elastic's 2026 investigation.

Campaign-specific observation

15,000+

Fake business profiles alleged in the 2026 Premium Home Service complaint.

Pending civil allegations

Unknown

The share of the combined market that remains undetected.

Why the denominator is missing

01 / WHAT THIS CASE STUDY INVESTIGATES

When online credibility becomes something an operator can manufacture

A hacked website can lend its established address to a page its owner never approved. A fabricated review can make a seller look trusted. A network of invented people can make a political message look like a public movement, the tactic usually called astroturfing.

These are connected problems of organic discovery and credibility, but they are not one documented conspiracy. This comparative case study follows the evidence across server intrusions, poisoned search results, review brokers, reputation suppression and covert political influence. Its central question is practical: **who controls the machinery, who benefits, and what do investigators actually know?**

"Organic" describes the route to an audience: ordinary search results, business listings, reviews, posts and replies. An operator may charge a client for manipulating those channels. That payment does not turn the resulting distribution into a paid platform placement. Advertising campaigns, impression fraud and advertising economics are outside this study.

Does this require infected computers?

Sometimes. A compromised web server is a computer; it can manipulate what a search engine sees without infecting a reader's laptop. In another chain, a poisoned search result leads the reader to run malware. Fake reviews and artificial political personas can operate without either infection. The case evidence below identifies which mechanism is actually documented.

Four mechanisms that should not be collapsed into “bots”

Mechanism	Asset under control	Organic objective	Does it require a hacked device?
Hacked-site SEO	Web server, site files or publishing environment	Inject pages or links; substitute crawler content; redirect visitors	Yes, for the intrusion cases here
Search-delivered malware	A lure and deceptive download	Reach people making relevant searches	The visitor is infected after a further action in the documented chains
Fake reviews and profiles	Invented identities and reputation signals	Influence local discovery and purchase decisions	No infection established in the review cases
Covert political influence	False personas, channels, content and coordination	Shape apparent agreement and message visibility	Optional; some operations also use hacking

The table is an editorial synthesis of the cited cases, not a claim that every campaign uses all four mechanisms.

02 / DEMONSTRATION

How SEO poisoning works: where hacking meets rankings



Original conceptual illustration. Two responses from one website do not necessarily mean the visitor’s device is infected.

mechanisms: finding them does not establish a particular ranking gain, sustained traffic increase or conversion rate.

The website is the first victim

A server intrusion gives the attacker a place inside a website’s delivery system. The attacker can then make the established domain publish something the owner did not intend. Google’s policies describe hacked pages, injected content (often called SEO spam) and malicious redirects; they also prohibit cloaking that deliberately shows different material to users and search engines for manipulation. A page can look ordinary to its administrator while a crawler or a selected visitor receives a different response. [Google Search spam policies](#).

A planted backlink tries to manufacture an endorsement. A new keyword page tries to capture a search. A substituted response tries to borrow an existing domain’s credibility. These are intended

Then the searcher can become the next victim

In a separate chain, a relevant-looking search result presents a supposed document or software installer. The user downloads and executes malicious code. Search visibility supplies the introduction; the execution supplies the endpoint infection. This distinction is essential to understanding Gootloader and the newer software-download campaigns.

CHAIN A · SERVER FIRST

- 1 Unauthorized access**
An exposed application is exploited or another foothold is obtained.
- 2 Persistent control**
A backdoor or server module lets the operator change responses.
- 3 Organic manipulation**
Injected pages, links or redirects influence discovery.

CHAIN B · SEARCH FIRST

- 1 Relevant search**
A person searches for a document or application.
- 2 Deceptive result**
The landing page offers a malicious download.
- 3 Endpoint infection**
Running the file gives the attacker access.


These diagrams summarize observed patterns; they are not a universal sequence. A fake review or coordinated social post does not prove either chain happened. Malware is also broader than a “virus”: server extensions, backdoors and trojanized installers need not reproduce like a traditional virus.

03 / RECENT SEO EXPOSURES

Black hat SEO examples: hacked websites, entry routes and missing evidence

The following cases prioritize 2026 disclosures and enforcement, with older investigations where they supply unusually clear commercial or infection evidence. A publication date, an observation period and a police action date are different events.

Gambling Goblin: legitimate domains, substituted pages, invented approval



Original conceptual illustration. A handful of compromised servers can push manipulated pages into ordinary search results; most results stay legitimate.

Check Point exposed compromised Brazilian servers, including government infrastructure, serving attacker-controlled pages through malicious Apache modules. Some pages imitated familiar app stores and displayed fabricated ratings and review counts. Researchers assessed search manipulation as a likely objective and had tracked activity since mid-2025. They did **not observe the original entry method**.

The important intersection is intrusion plus manufactured credibility: a legitimate domain carries the substituted page, while false ratings encourage trust. Those graphics were on counterfeit pages; this is not proof that reviews were posted to genuine app stores. The report identifies a Chinese-speaking actor, not publicly named defendants or a proven list of paying customers. It reports an exposure, not arrests. [Check Point's Gambling Goblin investigation](#).

UAT-10147: a documented answer to “how did they get in?”

Cisco Talos described exploitation of known weaknesses in internet-facing web applications, including Telerik and AjaxPro environments, plus an intrusion involving exposed ASP.NET cryptographic keys. After gaining code execution, operators installed tools that maintained control and supported SEO fraud. A recovered list of approximately 170,000 target URLs was a **target list, not 170,000 confirmed infections**. [Talos's initial-access investigation](#).

A companion analysis separates access tools from content-manipulation tools. BadIIS and a server-side handler supplied fabricated crawler content or scripts to selected visitors, while other components supported persistence and credential theft. The distinction explains why “ranking malware” is an incomplete description: the intrusion has multiple jobs. These reports do not identify ordinary retail customers or announce convictions. [Talos's UAT-10147 toolkit analysis](#).

BadIIS: a reusable capability, not one proven organization

Elastic reported more than 1,800 compromised Windows web servers in its investigated campaign. In a November 2025 intrusion, activity progressed from an observed web-shell foothold to the malicious IIS module in under 17 minutes. The original entry vector was unknown. The elapsed time describes that observed incident, not a universal time to hack a server. [Elastic's BADIIS investigation](#).

Talos later analyzed a builder supporting configurable redirects, crawler-facing content, metadata changes and link injection. It assessed with moderate confidence that the variant was privately sold or shared among groups. This supports a reusable tooling ecosystem, but not a verified sales ledger, current price or named customer list. [Talos's commodity BadIIS investigation](#).

Attribution needs similar care. Unit 42's Operation Rewrite, uncovered in March 2025, had infrastructure overlap with ESET Group 9. It reported **no infrastructure overlap with DragonRank**; similarities in functionality and tooling were a weaker connection. Shared software does not establish shared ownership. [Unit 42's Operation Rewrite analysis](#).

16 JUNE 2026 · INCIDENT INVESTIGATION

A fake WordPress plugin turns the footer into a backlink service

Sucuri found a plugin posing as a private-blog-network component. It retrieved outside content and placed it in a compromised WordPress site's footer; two web shells were stored in database records. This illustrates the practical link between unauthorized access and backlink manipulation: control of a site becomes the ability to place apparently independent endorsements.

The report documents an investigated intrusion, not a measured global network. It does not establish how the attacker originally entered or that the owner deliberately installed the malicious plugin. A backlink destination may benefit, but its presence alone does not identify a knowing customer. [Sucuri's WordPress PBN investigation](#).

A separate April 2026 Joomla case showed injected server-side code retrieving instructions for spam content or redirects. The operator could change what the site returned from outside the site itself. [Sucuri's Joomla investigation](#).

10 SEPTEMBER 2024 · FOUNDATIONAL COMMERCIAL EVIDENCE

DragonRank: customers could specify websites and keywords


Talos connected a Chinese-speaking operator cluster to more than 35 compromised IIS servers. Web-application exploitation led to web shells, credential theft and additional malware. Its public commercial presence offered customized search-promotion services and communication through messaging platforms. Customers could specify destination websites and keywords.

This is unusually direct evidence that hacking infrastructure can support an SEO service business. It does not disclose every customer or prove that every promised ranking result was delivered. "DragonRank" is a researcher's tracking label, not a person's legal name. The report does not announce arrests. [Talos's DragonRank investigation](#).

SEO poisoning: how a search result becomes an infected computer

ACTIVITY FROM 27 OCTOBER 2025 · INCIDENT EVIDENCE

Gootloader: a legitimate document search becomes execution



Original conceptual illustration. Search supplies the introduction; running the download supplies the infection.

Huntress documented a Bing search for a utility-easement document whose first result was a compromised website. A deceptive download led to malicious JavaScript execution and further access. The report explains the visitor's infection chain, but not the website's original break-in. A high search position and a familiar document topic provided credibility; they did not make the file safe. [Huntress's Gootloader investigation](#).

The causal boundary matters: a user searching and reading a result is different from executing the offered program. This case supplies evidence for search-delivered malware, not proof that the infected laptop subsequently wrote fake reviews or boosted another website.

Storm-2561: fake VPN software captures credentials

Microsoft described poisoned search results promoting fake VPN-client sites and trojanized installers. The counterfeit client could steal credentials, report an error and redirect the user toward legitimate software. Associated repositories were removed and certificates revoked. The announcement describes research and infrastructure action, not arrests. It does not show infected machines being repurposed into a ranking or political bot network. [Microsoft's Storm-2561 investigation](#).

OfferLoader / CL-CRI-1171: infection becomes a service

Unit 42 linked an SEO-poisoning funnel and gaming-video channels to a common loader serving a pay-per-install marketplace. Its investigated intrusions involved searches for legitimate utilities followed by trojanized downloads. Different malware buyers could receive delivery through the same loader, so one infected machine could host unrelated payloads.

Researchers identified more than 10,000 distinct loader samples. That counts files, not 10,000 unique victims, clients or operators. YouTube terminated eleven identified channels after notification. This is evidence of outsourced malware distribution through organic discovery, not an announced arrest or proof that every channel viewer became infected. The actors' legal identities and complete buyer list were not established. [Unit 42's infection-marketplace investigation](#).

SocGholish: a real takedown, with an important SEO limit

Dutch police reported that international partners remediated 14,971 infected WordPress websites and took down 106 servers and domains. Compromised sites presented fake updates; a visitor installing one could give attackers access to their computer. This establishes the website-to-endpoint infection route.

The announcement does not establish that poisoned rankings generated all those visits. Accordingly, it belongs here as supporting infection evidence, not as a proven takedown of 14,971 ranking-manipulation sites. The website count is not a count of cleaned personal computers, and the infrastructure total is not a number of arrested people. [Dutch police's SocGholish announcement](#).

Where the 3ve story fits, and where it stops

DOJ's 2018 announcement alleged access to more than 1.7 million malware-infected computers in the relevant botnet scheme. Its 2022 update identified Sergey Ovsyannikov and Yevgeniy Timchenko as arrested, extradited, guilty-pleading and sentenced participants; Aleksandr Isaev remained at large as of that release. This establishes infected-device infrastructure and real prosecution. It does **not** establish that those machines manipulated organic rankings, wrote reviews or posted political messages. The underlying business model is outside this case study's scope. Nor should Aleksandr Zhukov's separate Methbot sentence be assigned to these defendants. [DOJ, 2018](#); [DOJ, 2022](#).

The current organic cases provide their own evidence of entry routes. It would be speculative to transplant the infection details or objectives of one operation into another merely because both used malware.

05 / OPERATORS, BUYERS AND BENEFICIARIES

Who buys fake reviews and manufactured credibility?

There is no single customer profile. The evidence includes ordinary local businesses purchasing reviews, software businesses buying apparent approval, merchants using illicit marketplace consultants, reputation clients trying to bury reporting, and covert political operators seeking influence. The strongest identification comes from contracts, correspondence, complaints and court records. A destination link or a favorable political post, standing alone, is weaker evidence.

The operator

Controls the intrusion, review production, false identities or publishing infrastructure.

The customer

Commissions or pays for the service. A documented transaction is the clearest evidence.

The beneficiary

Receives traffic, favorable ratings or narrative support, sometimes without knowing or requesting it.

Reviews have a documented local-search connection. Google says local results mainly reflect relevance, distance and prominence; review volume and positive ratings can help local ranking. That does not mean every new review raises a website's general organic position, or that reviews override the other factors. App-store placement, Amazon product discovery, Maps results and ordinary web search must be analyzed separately. [Google's explanation of local ranking](#).

Who runs the services, and who pays?

The evidence describes a supply chain, not one kind of hacker. Intruders obtain control of websites; developers supply malicious tools; brokers sell visibility or reputation services; reviewers and account operators manufacture apparent endorsement. A customer may purchase a promised result without operating the machinery. In some cases, the same

business runs several stages. The clearest way to establish who benefited is to follow contracts, correspondence and court evidence, not simply identify the destination of a suspicious link.

This distinction matters for attribution. A compromised university or government website is ordinarily a victim. A business receiving traffic is not automatically a knowing customer. The examples below identify named buyers where records support that relationship and preserve the difference between allegations, settlements and criminal convictions.

Roomster and AppWinn: a documented buyer–supplier relationship



Original conceptual illustration. Operator, broker and beneficiary are separate roles; a business that benefits is not automatically the buyer.

Roomster provides unusually clear evidence connecting purchased reputation to ranking ambitions. In its 2022 complaint, the FTC alleged that Roomster, operated by **John Shriber and Roman Zaks**, bought more than **20,000 reviews** from **Jonathan Martinez, doing business as AppWinn**. The complaint describes thousands of fabricated accounts and correspondence linking review orders to searches for rooms and roommates. This identifies both the alleged supplier and the commercial beneficiary. It establishes an intention to manipulate discovery; it does not independently measure how much a particular order changed rankings. [FTC complaint, paragraphs 27–38](#)

The resulting settlement prohibited Roomster and its owners from buying or incentivizing reviews. The announced \$36.2 million judgment and \$10.9 million in penalties were largely suspended after a \$1.6 million payment because of inability to pay; those headline figures are not amounts collected. Martinez separately settled. The FTC subsequently recorded entry of the Roomster order on September 5, 2023. These were civil outcomes. [FTC settlement announcement](#); [FTC report, Appendix A, page 8](#)

Premium Home Service: 15,000 alleged fake Google business profiles

On May 11, 2026, federal authorities and Illinois sued **B.E.S.T. GDR, LLC, operating as Premium Home Service, and founder/CEO Yosef Bernath**. The complaint alleges more than **15,000 fake business profiles on Google Search and Maps**, fabricated favorable reviews and false local identities. According to authorities, local-looking phone numbers routed to overseas call centers; the company collected fees and passed jobs to third-party contractors while representing itself as the local service provider. Here the alleged commercial benefit was customer acquisition through deceptive local discovery. [DOJ announcement](#)

This case does not establish malware infection or that every contractor knew about the alleged scheme. As checked October 13, 2026, the FTC still listed the proceeding as **pending**, with no settlement or order entered. It is a civil complaint, not an arrest or conviction. [FTC case record](#)

Amazon and the BBB: a court order against a fake review broker

Amazon and the Better Business Bureau sued the operators of ReviewServiceUSA in 2024, alleging that fabricated customer accounts supplied positive reviews for product listings and business profiles. The announcement describes purchasing sellers and businesses without naming those customers. [Amazon's lawsuit announcement](#)

On June 22, 2026, plaintiffs' counsel reported a **permanent injunction**, attorneys' fees and an order requiring the defendant's domains to be disabled and transferred. This is an identifiable civil enforcement result, reported by the lawyers who obtained it. Their reference to more than 75 websites taken down concerns a separate earlier case; it should not be assigned to this judgment. [Davis Wright Tremaine announcement](#)

How accessible are these services to ordinary people?

Published investigations document several sales channels: public service websites, ordinary-looking SEO and reputation agencies, direct solicitation, social-media or messaging contacts, and referrals through intermediaries. DragonRank's researchers found commercial presentation and messaging channels; Amazon describes review brokers using outside social and messaging services. These observations establish discoverability, not reliability or a verified present-day price list. [Cisco Talos](#); [Amazon transparency report, page 21](#)

New York's **Operation Clean Turf** showed that small-business customers could encounter fake reviews through routine marketing inquiries. Investigators posing as a Brooklyn yogurt-shop owner contacted SEO companies; some offered fabricated reviews. The 2013 investigation identified public freelance recruitment and payments to writers of **\$1–\$10 per review**. That is historical labor compensation, not a current customer tariff. Named providers included eBoxed and XVIO; business beneficiaries included Laser Cosmetica and US Coachways. Nineteen companies agreed to stop the practices and pay more than \$350,000 in penalties. This was a settlement involving 19 companies, not 19 arrests. [New York attorney general](#)

Eliminalia: reputation laundering and suppressed search results



Original conceptual illustration. Suppression works by burying accurate reporting under manufactured content, not by deleting it.

Reputation manipulation also works by making accurate reporting harder to find. Forbidden Stories' 2023 investigation examined nearly 50,000 internal Eliminalia records and identified roughly 1,500 clients across 50 countries. It connected founder **Diego “Dídac” Sánchez**’s business to manufactured content and abusive removal claims. Reported contracts included more than €12,000 from Luis Carlos Orrantia and €110,000 from Pedro Miguel Haces Barba; the clients did not respond or declined comment. Eliminalia’s lawyers challenged the investigation and cited client confidentiality. Crucially, the consortium said allegations of **hacking could not be verified**. This supports a black-hat search-suppression case, not a proven malware operation. [Forbidden Stories investigation](#)

OCCRP’s reporting supplies a documented intermediary route: Adar Capital Partners and Diego Marynberg hired Eliminalia through reputation firm ReputationUp, with recorded payments exceeding €400,000. It also reported at least €100,000 from Area S.p.A.; Area said it authorized **legal methods only** and considered the targeted reporting inaccurate or unfair. Such responses matter: hiring reputation assistance does not prove a customer knowingly commissioned unlawful tactics. These are selected historical contracts, not representative market prices or proof every requested removal succeeded. [OCCRP and IRPI investigation](#)

Do people actually go to prison?

Yes, but the offense and outcome must be stated accurately. Some cases establish extortion or fraud incorporating review manipulation; others establish civil violations. A security-company exposure, account removal or server seizure is not a conviction.

William Stanley: 37 months, 2016. The SEO and reputation operator pleaded guilty to Hobbs Act extortion and received a federal prison sentence plus \$174,888 restitution. DOJ described threats to publish fabricated negative comments and reviews unless businesses paid. Generational Equity was a former client that became an extortion victim. This establishes punishment for coercive reputation abuse, not a rule that every search-engine policy violation is criminal. [DOJ sentencing announcement](#)

PromoSalento: nine-month sentence, Italy, 2018. TripAdvisor says it traced over 1,000 review attempts linked to an operator selling packages to hospitality businesses and supported the prosecution. The Criminal Court of Lecce imposed nine months' imprisonment and roughly €8,000 in costs and damages for reviews written under false identity. The platform's report establishes the sentence; it does not establish how much time was actually served. It does not publicly name the operator. [TripAdvisor transparency report, page 22](#)

Amazon insider-bribery case: prison and home confinement, 2023. Joseph Nilsen received 18 months in prison; Kristen Leccese received two years' probation including six months' home confinement. DOJ describes consultants connecting seller clients to bribed insiders, misuse of confidential access, and manipulated competitor reviews and listings. Their convictions involved broader conspiracy offenses, with tax misconduct also applying to Nilsen. The case shows how corrupt access, consulting services and reputation attacks can combine. It is not a conviction solely for purchasing reviews, and it does not establish a Google-ranking hack. [DOJ sentencing announcement](#)

06 / ORGANIC POLITICAL MANIPULATION

Astroturfing and bot farms: manufacturing the appearance of a public movement

Political influence: manufacturing the appearance of public support



Original conceptual illustration. An artificial crowd can contain automated, human-operated and hybrid accounts.

Covert political manipulation is documented. The strongest investigations identify organizations, funding relationships, account networks, stolen material, and deliberate impersonation. What they establish is usually more specific than “bots controlled an election.” An operation can manufacture apparent agreement, introduce a story into public discussion, or attack an opponent’s credibility without investigators being able to measure votes changed. Equally, a politician who benefits from a campaign is not automatically its customer.

The cases below concern deceptive identities, unpaid distribution, hacked accounts, and concealed sponsorship. “Organic” describes how content reaches an audience; the people running the operation may still be paid. A fabricated account may be operated by a person, software, or both. It does not necessarily require an infected consumer computer.

September 2026: a commercial platform targeting Malaysia

In its September 10, 2026 report, Anthropic attributed an election-manipulation platform to **BBS Bilisim Teknolojileri**, an Istanbul technology company. Investigators described approximately **1,000 fake X accounts**, a fabricated news outlet, and invented dossiers. This was influence sold as a commercial capability: a customer would not need to build the entire operation internally.

The activity included artificial engagement supporting Malaysia’s sitting prime minister. That does not establish he commissioned it. The operator also pursued a contract with Malaysia’s communications regulator; Anthropic found **no evidence that pursuit succeeded**. Its assessment placed the campaign in Breakout Category Two: multiple platforms, but no demonstrated spread into authentic communities. Figures in the millions shown by the operator’s dashboard were not independently verified. Anthropic banned the associated account and organization. BBS Teknoloji, whose name differs slightly from the entity in the report, publicly denied any knowledge of or involvement in the activity and said it was investigating whether its name or former associates had been misused. [The Edge Malaysia, September 15, 2026](#). Its report announces a disruption, not an arrest or conviction. [Source: Anthropic, GTG-84005 investigation](#).

August 2026: a Russia-linked campaign wrapped in a think tank

OpenAI disclosed a campaign promoting the **International Burke Institute**, which presented itself as an Israeli expert institution. OpenAI assessed that the operators very likely originated in Russia. They generated promotional posts and replies across several social platforms while attempting to conceal their origin. The website supported its apparent authority with copied academic work: **34 of 36 sampled articles** had been copied from elsewhere, sometimes with false attribution.

The campaign promoted a Russia-favorable sovereignty index. OpenAI did not identify a proven government commissioner or establish how several real people associated with the institute related to the Russian operators. Genuine engagement appeared limited, although associated Telegram channels commonly counted 10,000–20,000 subscribers. Subscribers are not a count of persuaded voters. OpenAI banned the identified ChatGPT accounts. This case illustrates how fabricated institutional credibility can support deceptive amplification, even when the measured audience remains modest. [Source: OpenAI, August 25, 2026.](#)

The RT bot farm: 968 accounts, not 968 arrests

On July 9, 2024, the U.S. Justice Department announced the seizure of **two domains** and searches involving **968 social-media accounts** in an alleged Russian government-operated bot farm. X suspended the remaining identified accounts. DOJ alleged that an unnamed RT deputy editor-in-chief organized the software's development and that an FSB officer subsequently ran a private intelligence organization with Kremlin approval and financial support.

The operation used AI-assisted fictitious personas, often appearing American, to distribute Russian-government narratives. A documented example impersonated a U.S. constituent replying to a political candidate with a Putin video about Ukraine. The alleged sponsor and intended beneficiary were governmental; this was not presented as a retail service for ordinary buyers. The announcement establishes infrastructure disruption and account searches. It does not announce arrests, prove that a particular person changed political beliefs, or establish infected household computers as a prerequisite. [Source: DOJ bot-farm disruption announcement.](#)

Doppelgänger: identifiable contractors and a state customer

Treasury identified **Ilya Gambashidze**, founder of the **Social Design Agency**, and **Nikolai Tupikin**, CEO and owner of **Structura**, as participants in a Russian presidential-administration-directed influence campaign. It described more than 60 websites impersonating legitimate news organizations, amplified through misleading social identities. The United States sanctioned the individuals and companies on March 20, 2024. [Source: Treasury.](#)

On September 4, DOJ announced seizure action involving 32 domains. Its affidavit alleged direction by the Russian Presidential Administration, particularly **Sergei Kiriyenko**, and also named **ANO Dialog**. The organic mechanism paired counterfeit news websites with personas posing as local citizens and circulating links. Goals included weakening support for Ukraine and influencing elections. These were allegations supporting seizures, not convictions. [Source: DOJ.](#)

Disruption did not mean disappearance: the EEAS's March 2026 report describes continuing activity during 2025, including Hebrew-speaking audiences. [Source: EEAS.](#)

Team Jorge: private pitches, hacked accounts, and unproven sales claims

The February 2023 Story Killers investigation identified **Tal Hanan** as “Jorge” and his brother **Zohar Hanan** as “Nick.” Undercover reporters observed demonstrations involving hacked messaging accounts and fake-account campaigns. They independently confirmed that a message sent during a demonstration from Kenyan political adviser Dennis Itumbi’s account reached its recipient.

Reporters gained access while posing as intermediaries for an African client. Hanan quoted €15 million for a presidential campaign and a €6 million minimum for the proposed short assignment. These were sales pitches, not verified payments or average market prices. His claimed record of 33 presidential campaigns, with 27 successes, did not prove he changed those outcomes. The actual payer for the demonstrated Kenyan activity remained unclear. Hanan denied wrongdoing.

The case documents an expensive service offered through private relationships and intermediaries. Its exposure is not itself a conviction, and a candidate’s apparent benefit does not identify the customer. [Source: OCCRP’s original undercover investigation.](#)

The Internet Research Agency: an influence operation with an SEO department

The February 16, 2018 indictment charged **13 Russian individuals and three companies**. DOJ alleged that **Yevgeniy Prigozhin** financed the **Internet Research Agency** through Concord companies. Its organizational description included graphics, IT, finance, and **search-engine optimization departments**. That is a concrete connection between SEO expertise and political influence, although it does not establish a particular hacked-website technique.

The alleged operation used fabricated or stolen identities, apparently American profiles, pages, and groups, and contact with unwitting Americans. DOJ expressly stated that the indictment did not allege knowing American participation or that the charged conduct changed the election result. [Source: DOJ indictment announcement.](#)

The legal outcome needs equal attention. On March 16, 2020, prosecutors sought dismissal with prejudice for **Concord Management and Consulting and Concord Catering**, citing difficulties imposing meaningful punishment, protection of investigative methods, and changes in usable evidence. Judge Dabney Friedrich granted dismissal. This was not a set of convictions against the indicted operators. [Source: filed dismissal motion; court reporting on the ruling.](#)

How much did these influence operations actually achieve?

A 2023 study linked U.S. respondents’ Twitter feeds with longitudinal surveys. Just **1% of sampled users accounted for 70% of exposure** to identified Russian influence accounts. Researchers found no meaningful relationship with changes in attitudes, polarization, or voting behavior. That finding is bounded by the study’s sample, platform, and measurements. It does not establish that every influence operation is ineffective, and its 70% figure is **not a measure of bot traffic on Twitter**. [Source: Eady and colleagues, Nature Communications.](#)

The reporting standard should therefore distinguish the operator’s ambition, its observable activity, its actual audience, and demonstrated effects. Counterfeit consensus is a real, documented product. Successful control of public opinion is a much larger claim that requires separate evidence.

What has been caught, and what cannot be counted

What percentage of Twitter is bots? Does the “70%” claim hold up?

Not as a verified, current claim about the whole platform. A real 70% finding exists, but its denominator is much narrower. Without the original source of the claim, its exact origin remains unknown.

Pew Research Center’s April 2018 report estimated that suspected automated accounts supplied **66% of sampled tweeted links to popular websites**. Excluding links to twitter.com raised that estimate to **70%**. It did not measure all Twitter accounts, all posts, visitors, impressions or network requests. It also did not establish that those automated accounts were all malicious. [Pew’s report and footnote 3](#).

The sample contained 1,156,837 tweets from 140,545 accounts listing English as their profile language, linking to 2,315 popular websites, between July 27 and September 11, 2017. Researchers collected data through Twitter’s public stream and used a calibrated Botometer threshold to identify suspected automation. This historical sample cannot describe the entire platform in 2026. [Pew’s methodology](#).

Three widely cited figures that measure different things		
Published finding	What was counted	Necessary limitation
Pew, 2018: 66%; 70% excluding twitter.com	Sampled tweeted links to popular sites, using 2017 data.	Link output, not current platform-wide traffic.
Varol and colleagues, 2017: 9%–15%	Estimated share of active Twitter accounts classified as bots.	Historical, model-based account estimate; not a 2026 measurement.
Thales, 2026: 53% automated; 40% bad bots	Broader web traffic in its analysis of 2025 activity.	Vendor telemetry, not an X-specific census. Both percentages use total traffic as their denominator.

The Thales statistic belongs here only to clarify the claim. It includes applications, APIs and automation outside organic search or political messaging. A bot making thousands of requests is not thousands of people.

“Bot,” “fake” and “coordinated” are different findings

A bot automates activity. A fake persona misrepresents an identity. A coordinated influence network conceals relationships or sponsorship. A compromised account has been taken over without authorization. These categories can overlap, but none automatically proves the others. Human employees can operate false personas; legitimate services can post automatically.

Classification also carries uncertainty. Pew acknowledges false positives and false negatives and explains that its threshold served an aggregate estimate, rather than definitive judgments about individual accounts. [Pew’s explanation of detection limits](#). A separate study applying three detection methods to the same political-discussion dataset found substantially different classifications. [Martini and colleagues, 2021](#).

What has been caught, and what remains unknown?

Official disclosures show large enforcement volumes. They do not provide a combined detection rate for hacked SEO, purchased reviews and political influence operations. The following are platform-reported actions, not counts of arrests or convictions.

Google Maps: fake Google reviews and business identities

On April 16, 2026, Google reported that during 2025 it blocked or removed **more than 292 million policy-violating reviews**, removed **over 13 million fake Business Profiles**, restricted posting by more than 782,000 policy-violating accounts, and blocked 79 million inaccurate or unverified edits. It also reported publishing more than one billion helpful reviews. [Google's 2025 enforcement summary](#).

“Policy-violating” is broader than purchased or computer-generated reviews. “Blocked or removed” combines material stopped before publication with material taken down afterward. Dividing removals by published reviews would not reveal the proportion of visible reviews that are fake, or how many undetected frauds remain.

X: suspended accounts are not unique operators

In its April 24, 2026 submission to a UK parliamentary committee, X reported **282.6 million global account suspensions** under “Multiple Accounts and Coordination” for October 1, 2024–September 30, 2025. Its preceding annual period recorded 465.9 million under the differently named “Platform Manipulation and Spam” category. [X's parliamentary response, page 3](#).

Those categories do not isolate political bots, and their different labels prevent a straightforward trend claim. An annual flow of suspended accounts cannot be divided by a snapshot of active users to establish prevalence. Operators may repeatedly create disposable replacement accounts.

Meta: proactive detection does not mean almost everything was caught

Meta's May 2026 parliamentary update reported action against approximately **1.1 billion fake Facebook accounts in Q4 2025**; 99.9% were detected before a user reported them. Separately, it described six coordinated inauthentic behavior networks in its H1 2026 threat report, with approximately 1,450 Facebook accounts, 199 Pages and 704 Instagram accounts removed. [Meta's parliamentary update](#).

The 99.9% figure concerns accounts that received enforcement action. It does not mean Meta caught 99.9% of all fake accounts that existed. Nor does the billion-account total represent a billion people or political operatives.

The missing denominator



Original conceptual illustration. Platforms report what they caught; nobody can count what they missed.

No defensible universal “percentage still unnoticed” emerges from these disclosures. If investigators discover ten networks, the proportion caught depends on whether twelve, a hundred or ten thousand existed. The discovered total alone cannot answer that question.

Higher removal counts may reflect more attempted abuse, better detection, changed definitions, or several changes together. Lower counts can be equally ambiguous. Hidden operators, incomplete cross-platform coverage and overlapping investigations further complicate any worldwide total.

For political messaging, distinguish **production, delivery, human exposure, belief change and behavior change**. These are successive evidentiary steps. Posting a million messages does not establish a million human readers, and reaching readers does not establish changed votes. That distinction supports serious investigation without mistaking an operator’s ambition (or sales pitch) for a demonstrated result.

“Caught” describes several different outcomes

The case record supports real criminal punishment, civil orders and technical disruptions. It does not support calling every disclosed operation a prosecution. The examples below are cross-references to the detailed cases, not additional incidents to add to a global total.

How to read the legal and enforcement labels		
Outcome	What it establishes	Example in this study
Research exposure	Investigators identified activity or infrastructure within their visibility.	Gambling Goblin
Platform action	A service removed content, accounts or access under its rules.	2026 AI-provider influence disclosures
Police disruption	Specific infrastructure or access was interrupted.	SocGholish action
Civil complaint	Plaintiffs filed allegations; liability remains unresolved unless later adjudicated.	Premium Home Service
Settlement or injunction	A civil agreement or court order imposes obligations; its wording governs any admission or finding.	Roomster and Amazon/BBB
Indictment	Criminal charges were filed; it is not a conviction.	IRA defendants
Criminal sentence	A court imposed punishment after a plea or conviction.	Stanley, PromoSalento and Nilsen



Original conceptual illustration. Research exposure, platform action, disruption, complaint, settlement, indictment and sentence are different outcomes.

“Black hat” is an industry description, not a single criminal charge. Search-policy violations, deception, unauthorized access, extortion and bribery are different conduct. Their legal treatment depends on facts and jurisdiction. The named sentences in this study attach to the offenses described by the sources; they do not make ordinary SEO a crime. Equally, removing a bot

account is not the same as identifying or arresting the person behind it.

09 / IMPORTANT TAKEAWAYS

What the evidence allows us to conclude

1. **Hacking and black hat SEO do intersect.** Server intrusions can supply injected links, substitute indexed content or route search visitors elsewhere. Separately, poisoned search results can deliver malware.
2. **The business is accessible through intermediaries.** Public service pitches, SEO firms, reputation agencies and marketplace consultants appear in the evidence. That does not prove a buyer knew which hidden methods would be used.
3. **Reviews connect discovery with persuasion.** Some schemes seek placement; others seek to convert the person who already found a listing. Ranking intent and measured ranking effects are different findings.
4. **Political influence services exist.** False identities and coordinated distribution can be purchased or state-sponsored. An apparent beneficiary is not automatically the customer.
5. **People have been sentenced.** The strongest examples concern identified offenses such as extortion, identity deception, bribery or broader fraud. Many recent SEO exposures have no publicly announced arrest.
6. **Large counts need narrow labels.** Websites, accounts, reviews, file samples, posts and people are not interchangeable units.
7. **The hidden remainder is not reliably quantified.** This research establishes substantial observed activity, not the fraction of the worldwide market that escapes detection.

A reusable checklist for the next investigation

Record the operator's identity or tracking label; the customer evidence; the benefiting destination; the compromised asset, if any; the first observed access; the organic distribution channel; the exact date and unit of every count; the source's attribution confidence; the latest legal status; and what remains unproved. Preserve denials and distinguish a claimed capability from an observed result.

The most useful follow-up evidence would connect these layers: a contract showing what a customer ordered, technical records showing what an operator actually deployed, and independent measurements showing what real people encountered. Without those links, a story may accurately expose machinery while remaining unable to identify every buyer or establish its effect.

10 / FREQUENTLY ASKED QUESTIONS

Black hat SEO, fake reviews and bot farms: common questions

What is black hat SEO?

Black hat SEO is any attempt to rank or look trustworthy in search by breaking search engine rules or deceiving people. In this study that includes hacked websites carrying injected pages and links, cloaking that shows search engines different content than visitors, purchased fake reviews and fake business profiles, and campaigns that bury accurate reporting. It is an industry label, not a single crime: the legal risk depends on the specific conduct.

Does black hat SEO still work?

It is still sold and still used. Researchers documented DragonRank offering ranking services built on hacked servers, and the 2026 BadIIS and Gambling Goblin campaigns show the tooling is active. What the evidence rarely shows is a measured, lasting ranking gain. Google's spam policies target hacked content, cloaking and sneaky redirects, Google reported removing more than 13 million fake Business Profiles in 2025, and operators and buyers in this study have faced FTC orders, lawsuits and prison sentences.

What is SEO poisoning?

SEO poisoning means pushing malicious or deceptive pages into ordinary search results so that people searching for a document or a piece of software land on them. In the Gootloader case, a search for a legal document led to a compromised site and a malicious download. In Storm-2561, poisoned results promoted fake VPN installers that stole credentials. The search result supplies the trust; running the download supplies the infection.

How can I tell if my website was hacked for SEO spam?

Common signs include pages in Google's index that you never published, unfamiliar links in your footer, content that looks different to Googlebot than it does to you, redirects that only affect some visitors, and unknown plugins or database entries. Google Search Console's Security Issues report flags many of these. The Sucuri WordPress and Joomla cases in this study show how a fake plugin and injected server code produced exactly these symptoms.

Are fake Google reviews legal?

In the United States, no. The FTC finalized a rule in August 2024 that bans fake reviews, including AI-generated ones, and bans paying for reviews that must be positive or negative; knowing violators can face civil penalties. [FTC final rule announcement](#). Fake reviews also break Google's policies. Enforcement in this study ranges from the FTC's Roomster order to the pending Premium Home Service complaint and a prison sentence in Italy's PromoSalento case. This is general information, not legal advice.

Can fake Google reviews be traced?

Sometimes, and usually through records rather than the reviews alone. The FTC’s Roomster complaint relied on correspondence tying more than 20,000 purchased reviews to a named supplier. TripAdvisor traced over 1,000 review attempts to one operator, which supported a criminal case. Platforms also catch patterns at scale: Google says it blocked or removed more than 292 million policy-violating reviews in 2025.

What is an example of astroturfing?

Astroturfing is fake grassroots support. The Doppelgänger campaign paired counterfeit news websites with personas posing as local citizens. The U.S. Justice Department’s 2024 action against an alleged Russian government bot farm involved 968 accounts, many posing as Americans. In 2026, Anthropic described a commercial platform using about 1,000 fake X accounts to boost political messaging in Malaysia; the company it named has denied involvement.

What is a bot farm, and is it illegal?

A bot farm is a coordinated set of automated or fake accounts used to amplify messages, inflate engagement or impersonate real people. Automation by itself is not a crime. Legal exposure comes from what the farm is used for: the 2024 DOJ seizures targeted an alleged foreign government influence operation, and the FTC’s 2024 rule bans buying or selling fake followers or views when the buyer knew or should have known they were fake. Platforms also suspend these accounts by the hundreds of millions.

What percentage of Twitter (X) is bots?

Nobody has a verified current figure. The often-repeated 70% comes from a 2018 Pew study that measured links to popular websites shared by suspected bots, not accounts or traffic. A 2017 academic estimate put bots at 9% to 15% of active accounts. X reports suspensions, not prevalence, and suspension counts cannot be turned into a share of active users.

11 / REFERENCE NOTES

Method and reuse

Research cutoff: 23 September 2026; legal statuses and company responses refreshed 13 October 2026. This is a curated investigation of recent public disclosures and strong historical precedents, not an exhaustive worldwide incident census. Primary technical investigations, official court/enforcement releases, complaints, platform disclosures and original investigative reporting underpin the case summaries. Statements about “pending” proceedings or undisclosed customers are limited to the reviewed public record at this cutoff.

Security vendors see their own telemetry and investigations. Platforms report their own enforcement. Plaintiffs and their counsel describe their cases from a litigant’s position. Journalistic investigations can reveal contracts and interviews that technical telemetry cannot. These sources answer different questions; their counts and certainty levels should travel with any quotation or paraphrase.

Illustrations are explanatory artwork generated for this study, not documentary photographs, recovered evidence or screenshots. The delivery diagrams are editorial syntheses. The writing separates established technical observations, allegations, reported court outcomes and analysis. The 70% fact check preserves the historical sample rather than turning it into a current estimate.

This is a comparative investigation, not a claim that its separate campaigns had one operator. For a single-campaign deep dive in the same format, see the [Optimization Theory Gootloader SEO poisoning case study](#).

If you quote or reuse this research, keep the source link beside each case and check the latest status of pending proceedings. Do not turn a dated allegation into an established fact. The source index below collects every reference cited in the narrative.

12 / SOURCE INDEX

Follow the evidence

54 linked references, checked for this research on 23 September 2026 and refreshed 13 October 2026. Links are also placed beside the claims they support. Search by publisher or title; the full article remains visible above.

01

[Google Search spam policies](https://developers.google.com/search/docs/essentials/spam-policies)
· <https://developers.google.com/search/docs/essentials/spam-policies>
developers.google.com

02

[Check Point’s Gambling Goblin investigation](https://research.checkpoint.com/2026/gaming-the-system-how-a-chinese-speaking-actor-turned-brazilian-government-sites-into-an-seo-weapon/)
· <https://research.checkpoint.com/2026/gaming-the-system-how-a-chinese-speaking-actor-turned-brazilian-government-sites-into-an-seo-weapon/>
research.checkpoint.com

03

[Talos’s initial-access investigation](https://blog.talosintelligence.com/uat-10147-chinese-speaking-adversary-integrates-agentic-ai-into-post-compromise-operations/)
· <https://blog.talosintelligence.com/uat-10147-chinese-speaking-adversary-integrates-agentic-ai-into-post-compromise-operations/>
blog.talosintelligence.com

04

[Talos’s UAT-10147 toolkit analysis](https://blog.talosintelligence.com/uat-10147-deploys-spectre-a-cross-platform-implant-with-linux-rootkit-and-byovd-capabilities/)
· <https://blog.talosintelligence.com/uat-10147-deploys-spectre-a-cross-platform-implant-with-linux-rootkit-and-byovd-capabilities/>
blog.talosintelligence.com

05

[Elastic’s BADIIS investigation](https://security-labs.elastic.co/security-labs/badiis-to-the-bone-new-insights-to-global-seo-poisoning-campaign)
· <https://security-labs.elastic.co/security-labs/badiis-to-the-bone-new-insights-to-global-seo-poisoning-campaign>
security-labs.elastic.co

-
- 06 [Talos's commodity BadIIS investigation](#)
· <https://blog.talosintelligence.com/from-pdb-strings-to-maas-tracking-a-commodity-badiis-ecosystem/>
blog.talosintelligence.com
-
- 07 [Unit 42's Operation Rewrite analysis](#)
· <https://unit42.paloaltonetworks.com/operation-rewrite-seo-poisoning-campaign/>
unit42.paloaltonetworks.com
-
- 08 [Sucuri's WordPress PBN investigation](#)
· <https://blog.sucuri.net/2026/06/wordpress-pbn-plugin-drops-dual-webshells-via-database-injection.html>
blog.sucuri.net
-
- 09 [Sucuri's Joomla investigation](#)
· <https://blog.sucuri.net/2026/04/joomla-seo-spam-injector-obfuscated-php-backdoor-hijacking-site-visitors.html>
blog.sucuri.net
-
- 10 [Talos's DragonRank investigation](#)
· <https://blog.talosintelligence.com/dragon-rank-seo-poisoning/>
blog.talosintelligence.com
-
- 11 [Huntress's Gootloader investigation](#)
· <https://www.huntress.com/blog/gootloader-threat-detection-woff2-obfuscation>
www.huntress.com
-
- 12 [Microsoft's Storm-2561 investigation](#)
· <https://www.microsoft.com/en-us/security/blog/2026/03/12/storm-2561-uses-seo-poisoning-to-distribute-fake-vpn-clients-for-credential-theft/>
www.microsoft.com
-
- 13 [Unit 42's infection-marketplace investigation](#)
· <https://unit42.paloaltonetworks.com/ppi-network-malware-campaign-analysis/>
unit42.paloaltonetworks.com
-
- 14 [Dutch police's SocGholish announcement](#)
· <https://www.politie.nl/en/news/2026/juni/18/11-international-law-enforcement-initiate-hunt-on-malware-group-socgholish.html>
www.politie.nl
-
- 15 [DOJ, 2018](#)
· <https://www.justice.gov/usao-edny/pr/two-international-cybercriminal-rings-dismantled-and-eight-defendants-indicted-causing>
www.justice.gov
-
- 16 [DOJ, 2022](#)
· <https://www.justice.gov/usao-edny/pr/united-states-recovers-over-15-million-swiss-bank-accounts-proceeds-global-digital>
-

-
- 17 [Google's explanation of local ranking](https://support.google.com/business/answer/7091?hl=en)
· <https://support.google.com/business/answer/7091?hl=en>
support.google.com
-
- 18 [FTC complaint, paragraphs 27–38](https://www.ftc.gov/system/files/ftc_gov/pdf/Roomster%20Complaint.pdf)
· https://www.ftc.gov/system/files/ftc_gov/pdf/Roomster%20Complaint.pdf
www.ftc.gov
-
- 19 [FTC settlement announcement](https://www.ftc.gov/news-events/news/press-releases/2023/08/ftc-state-partners-secure-proposed-order-banning-roomster-owners-using-deceptive-reviews)
· <https://www.ftc.gov/news-events/news/press-releases/2023/08/ftc-state-partners-secure-proposed-order-banning-roomster-owners-using-deceptive-reviews>
www.ftc.gov
-
- 20 [FTC report, Appendix A, page 8](https://search.ftc.gov/system/files/ftc_gov/pdf/p238400_ftc_collaboration_act_report.pdf)
· https://search.ftc.gov/system/files/ftc_gov/pdf/p238400_ftc_collaboration_act_report.pdf
search.ftc.gov
-
- 21 [DOJ announcement](https://www.justice.gov/opa/pr/departments-justice-files-complaint-against-best-gdr-llc-doing-business-premium-home-service)
· <https://www.justice.gov/opa/pr/departments-justice-files-complaint-against-best-gdr-llc-doing-business-premium-home-service>
www.justice.gov
-
- 22 [FTC case record](https://www.ftc.gov/legal-library/browse/cases-proceedings/best-gdr-llc-et-al-united-states-state-illinois-v)
· <https://www.ftc.gov/legal-library/browse/cases-proceedings/best-gdr-llc-et-al-united-states-state-illinois-v>
www.ftc.gov
-
- 23 [Amazon's lawsuit announcement](https://www.aboutamazon.com/news/policy-news-views/amazon-better-business-bureau-file-lawsuit-on-fake-reviews)
· <https://www.aboutamazon.com/news/policy-news-views/amazon-better-business-bureau-file-lawsuit-on-fake-reviews>
www.aboutamazon.com
-
- 24 [Davis Wright Tremaine announcement](https://www.dwt.com/about/news/2026/06/dwt-wins-fake-review-case-amazon-bbb)
· <https://www.dwt.com/about/news/2026/06/dwt-wins-fake-review-case-amazon-bbb>
www.dwt.com
-
- 25 [Amazon transparency report, page 21](https://assets.aboutamazon.com/cd/28/4d02dd2e41ec8c6d1bc341e9d919/amazon-eu-store-transparency-report-jan-june-2023.pdf)
· <https://assets.aboutamazon.com/cd/28/4d02dd2e41ec8c6d1bc341e9d919/amazon-eu-store-transparency-report-jan-june-2023.pdf>
assets.aboutamazon.com
-
- 26 [New York attorney general](https://ag.ny.gov/press-release/2013/ag-schneiderman-announces-agreement-19-companies-stop-writing-fake-online-reviews)
· <https://ag.ny.gov/press-release/2013/ag-schneiderman-announces-agreement-19-companies-stop-writing-fake-online-reviews>
ag.ny.gov

-
- 27 [Forbidden Stories investigation](#)
· <https://forbiddenstories.org/the-gravediggers-eliminialia/>
forbiddenstories.org
-
- 28 [OCCRP and IRPI investigation](#)
· <https://www.occrp.org/en/project/story-killers/eliminialia-a-reputation-laundromat-for-criminals>
www.occrp.org
-
- 29 [DOJ sentencing announcement](#)
· <https://www.justice.gov/usao-ndtx/pr/search-engine-optimizer-sentenced-more-three-years-federal-prison-extorting-money-local>
www.justice.gov
-
- 30 [TripAdvisor transparency report, page 22](#)
· https://www.tripadvisor.com/TripAdvisorInsights/wp-content/uploads/2019/09/2147_PR_Content_Transparency_Report_6SEP19_US.pdf
www.tripadvisor.com
-
- 31 [DOJ sentencing announcement](#)
· <https://www.justice.gov/usao-wdwa/pr/final-two-us-based-defendants-amazon-bribery-case-sentenced>
www.justice.gov
-
- 32 [Source: Anthropic, GTG-84005 investigation](#)
· <https://www.anthropic.com/threat-intelligence-report-september-2026>
www.anthropic.com
-
- 33 [Source: OpenAI, August 25, 2026](#)
· <https://openai.com/index/disrupting-malicious-uses-of-ai-influence-campaign-russia/>
openai.com
-
- 34 [Source: DOJ bot-farm disruption announcement](#)
· <https://www.justice.gov/archives/opa/pr/justice-department-leads-efforts-among-federal-international-and-private-sector-partners>
www.justice.gov
-
- 35 [Source: Treasury](#)
· <https://home.treasury.gov/news/press-releases/jy2195>
home.treasury.gov
-
- 36 [Source: DOJ](#)
· <https://www.justice.gov/usao-edpa/pr/justice-department-disrupts-covert-russian-government-sponsored-foreign-malign>
www.justice.gov
-
- 37 [Source: EEAS](#)
· <https://euvsdisinfo.eu/eeas-4th-fimi-threat-report-march-2026/>

-
- 38 [Source: OCCRP's original undercover investigation](#)
· <https://www.occrp.org/en/project/story-killers/hacks-bots-and-blackmail-how-secret-cyber-mercenaries-disrupt-elections>
www.occrp.org
-
- 39 [Source: DOJ indictment announcement](#)
· <https://www.justice.gov/archives/opa/pr/grand-jury-indicts-thirteen-russian-individuals-and-three-russian-companies-scheme-interfere>
www.justice.gov
-
- 40 [Source: filed dismissal motion](#)
· <https://www.courthousenews.com/wp-content/uploads/2020/03/ConcordMgmt-USDismissal.pdf>
www.courthousenews.com
-
- 41 [court reporting on the ruling](#)
· <https://www.courthousenews.com/feds-move-to-drop-charges-against-russian-firm-filed-in-mueller-probe/>
www.courthousenews.com
-
- 42 [Source: Eady and colleagues, Nature Communications](#)
· <https://www.nature.com/articles/s41467-022-35576-9>
www.nature.com
-
- 43 [Pew's report and footnote 3](#)
· <https://www.pewresearch.org/internet/2018/04/09/bots-in-the-tweetsphere/>
www.pewresearch.org
-
- 44 [Pew's methodology](#)
· <https://www.pewresearch.org/internet/2018/04/09/bots-in-the-tweetsphere-methodology/>
www.pewresearch.org
-
- 45 [Varol and colleagues, 2017: 9%–15%](#)
· <https://arxiv.org/abs/1703.03107>
arxiv.org
-
- 46 [Thales, 2026: 53% automated; 40% bad bots](#)
· <https://cpl.thalesgroup.com/blog/application-security/bad-bots-in-the-agentic-age>
cpl.thalesgroup.com
-
- 47 [Pew's explanation of detection limits](#)
· <https://www.pewresearch.org/short-reads/2018/04/19/qa-how-pew-research-center-identified-bots-on-twitter/>
www.pewresearch.org

48

[Martini and colleagues, 2021](#)

[· https://journals.sagepub.com/doi/full/10.1177/20539517211033566](https://journals.sagepub.com/doi/full/10.1177/20539517211033566)

journals.sagepub.com

49

[Google’s 2025 enforcement summary](#)

[· https://blog.google/products-and-platforms/products/maps/new-ways-were-protecting-businesses-on-maps/](https://blog.google/products-and-platforms/products/maps/new-ways-were-protecting-businesses-on-maps/)

blog.google

50

[X’s parliamentary response, page 3](#)

[· https://committees.parliament.uk/publications/53682/documents/299592/default/](https://committees.parliament.uk/publications/53682/documents/299592/default/)

committees.parliament.uk

51

[Meta’s parliamentary update](#)

[· https://committees.parliament.uk/publications/53679/documents/299534/default/](https://committees.parliament.uk/publications/53679/documents/299534/default/)

committees.parliament.uk

52

[Optimization Theory Gootloader case study](#)

[· https://optimizationtheory.com/project-view/seo-poisoning-gootloader/](https://optimizationtheory.com/project-view/seo-poisoning-gootloader/)

optimizationtheory.com

53

[The Edge Malaysia, BBS Teknologi denial](#)

[· https://theedgemalaysia.com/node/818221](https://theedgemalaysia.com/node/818221)

theedgemalaysia.com

54

[FTC final rule banning fake reviews](#)

[· https://www.ftc.gov/news-events/news/press-releases/2024/08/federal-trade-commission-announces-final-rule-banning-fake-reviews-testimonials](https://www.ftc.gov/news-events/news/press-releases/2024/08/federal-trade-commission-announces-final-rule-banning-fake-reviews-testimonials)

www.ftc.gov