Technical and Structural Mechanisms of Infrastructure-Level Censorship: Assessing Network-Aligned Interventions in Content Delivery, DNS Registrar Systems, and Automated Indexing Engine Pipelines

Technical Taxonomy of Network Interventions

Digital content moderation has evolved from reactive, human-reviewed content removal into a proactive, algorithmic, and infrastructure-level system of containment. This structural shift is driven by a network of corporate, philanthropic, and state-aligned entities that operate across multiple layers of global internet architecture. By intervening at the Domain Name System (DNS), Content Delivery Network (CDN), and Application Programming Interface (API) levels, these entities systematically bypass the legal and constitutional barriers that protect independent journalism and open-source intelligence (OSINT) investigations.

To understand how this technical apparatus operates, it is necessary to trace the flow of capital and software development that fuels these interventions. At the core of the funding infrastructure is the Vine & Fig Tree network, which utilizes a tripartite corporate architecture to direct capital into software engineering and custom linguistic datasets. This network includes:

- **Vine & Fig Tree Institute l Inc. (501(c)(3))**: Serves as a primary tax-exempt capital clearinghouse, securing over $3 million in its initial fiscal year ending December 2024 to function as a financial conduit.
- **Vine & Fig Tree Fund Inc. (501(c)(3) private foundation)**: Directs downstream grant allocations, transferring $850,000 to the Adir Challenge Foundation in Hoboken, New Jersey.
- **Vine & Fig Tree Action Inc. (501(c)(4) social welfare organization)**: Conducts lobbying, narrative mobilization, and legislative advocacy to influence policy outside the strict constraints of public charities.

The Adir Challenge Foundation, led by Morielle Lotan and former Israeli intelligence officer Dr. Shay Hershkovitz, serves as a rapid-prototyping incubator that crowdsources and scales closed-source software designed to automate digital discourse control. This incubator has produced specialized tools, including CTRL (a gaming plugin that calculates player "Vibe Scores" using sentiment analysis and Player Reputation APIs), Flaggy (a context-aware NLP model designed to parse sentence structure to detect verbal abuse), and The Unifiers (a platform that gamifies moderation by rewarding users for submitting detailed harassment reports). Additionally, platforms like Reportify utilize generative AI to convert unstructured descriptions of incidents into standardized Title VI civil rights legal complaints. This tool submits reports to multiple regulatory and educational portals simultaneously, systematically escalating administrative and legal pressure to force preemptive restriction of speech. Concurrently, Oct7

Community OS, led by former IDF technology commander Omer Dagan and Vintage Investment Partners founder Alan Feld, aggregates and synchronizes digital assets into automated, real-time narrative campaigns across mainstream social networks.
This custom-built software stack is validated and integrated into mainstream corporate technology infrastructures through strategic partnerships. Google Jigsaw's Head of Operations and Chief of Staff, Raquel Saxe, serves as an advisor to the Adir Challenge Foundation, establishing an administrative bridge to transition these AI models into Jigsaw's moderation architectures. By validating crowdsourced NLP models (like Flaggy) against Jigsaw's enterprise-grade Perspective API, the network embeds custom toxicity parameters directly into the automated filtering pipelines of major digital publishers and social networks.
Furthermore, legacy organizations—such as CyberWell, the Anti-Defamation League (ADL), and the USC Shoah Foundation (which runs the Countering Antisemitism Laboratory)—maintain "Trusted Flagger" or "Trusted Partner" status with major platforms like Meta, TikTok, and YouTube. This status provides them with direct API access to platform backends, allowing them to upload white-label threat-intelligence feeds and structured linguistic datasets that bypass standard manual moderation queues, resulting in the rapid deletion of content and account suspensions.

Infrastructure Layer	Primary Technical Vector	Aligned Operational Entities	Target Assets & Platforms	Systemic Moderation Outcome
DNS / Registrar Layer	Localized administrative codes, WHOIS audits, Terms-of-Service violations, DNS-sinkholing	Registry compliance monitors, network-aligned legal groups	Domain registrars, Top-Level Domain (TLD) registries, DNS zone files	Absolute domain revocation, traffic redirection, and total loss of domain resolution
CDN / Cloud Hosting Layer	Trust-and-safety policy interventions, legal liability notices, executive coordination	Combat Antisemitism Movement, Louis D. Brandeis Center	Cloudflare, Akamai, AWS edge nodes and origin hosting	De-platforming, loss of DDoS mitigation, exposure of origin servers to cyberattacks
API / Platform Moderation Layer	"Trusted Flagger" database synchronization, direct programmatic API moderation bypass	CyberWell, ADL, USC Shoah Foundation, Google Jigsaw	Meta, TikTok, YouTube content moderation queues, Google Jigsaw's Perspective API	Automated content deletion, instant account suspension, validation of custom filtering algorithms

CDN and Infrastructure Pressure Models

For alternative media repositories and independent journalism platforms that survive initial platform-level censorship, content delivery and hosting infrastructure represent the next major choke point. CDNs such as Cloudflare and Akamai, alongside cloud providers like Amazon Web Services (AWS), provide the caching, bandwidth, and distributed denial-of-service (DDoS) protection necessary to keep high-traffic alternative media sites online.

Network-aligned entities, including the Combat Antisemitism Movement (CAM) and the Louis D. Brandeis Center for Human Rights Under Law, deploy structured legal liability notices and trust-and-safety interventions directly targeting these infrastructure giants. CAM operates as an administrative hub funded through Adam Beren's Beren Sea Foundation, which has injected over $13.1 million into allied initiatives since 2020. It funds entities like the Network Contagion Research Institute (NCRI) to produce academic studies targeting online hate and the algorithmic amplification of conspiracy theories. These studies are then packaged into formal pressure models to force infrastructure providers to act.

The pressure model relies on presenting cloud providers with detailed dossiers that classify independent investigative journalism as "hate speech," "extremism," or "misinformation". By exposing these infrastructure giants to secondary liability risks and potential brand damage, the network forces trust-and-safety departments to enforce terms-of-service (ToS) clauses that prohibit the hosting of "harmful" or "unverified" document corpuses. A historical precursor to this model was documented when trust-and-safety compliance campaigns targeted Cloudflare to de-platform French-language websites posting controversial content, demonstrating that infrastructure providers are highly vulnerable to campaigns that link their services to the enablement of hate speech.

This pressure is reinforced by structural connections within the technology sector. For instance, Tom Leighton, the CEO and co-founder of Akamai Technologies and a professor of applied mathematics at MIT, has maintained close philanthropic and academic ties to institutions aligned with these centers, establishing named chairs like the Berger-Leighton Endowed Professorship at Brandeis University and supporting the Lavine Family Professor and Director of the Brandeis Center for Jewish Studies. These cross-institutional relationships align corporate leadership with the strategic goals of the network, making infrastructure providers highly receptive to formal outreach and trust-and-safety interventions. Risk-averse infrastructure monopolies frequently respond to these campaigns by terminating service agreements, exposing target sites to devastating cyberattacks and leaving alternative media repositories functionally offline.

Domain Registrar and DNS Level Targeting

When alternative media platforms attempt to circumvent CDN-level de-platforming by utilizing decentralized hosting, the network shifts its targeting to the baseline routing layer of the internet: domain registrars and the Domain Name System (DNS). This technical vector represents an elite form of censorship that avoids public debates over freedom of speech by focusing strictly on administrative and technical compliance.

Network-aligned legal teams systematically audit the registry information of target sites, scanning for minor administrative errors, WHOIS inaccuracies, or localized terms-of-service violations. By exploiting administrative loopholes and compliance clauses, these actors file formal complaints with domain registrars to secure the systemic de-registration of target domains.

If the registrar fails to act, the network leverages localized administrative rules to execute DNS sinkholing. Under standard routing configurations, clients attempting to resolve the target domain are redirected to a controlled, dummy IP address managed by a security operator or registry compliance authority. This redirection breaks the resolution chain, preventing users from accessing independent document corpuses, OSINT investigation databases, or UAP archives. Furthermore, security platforms employ passive DNS traffic analysis to monitor real-world DNS flows, identifying alternative domains registered by the same investigative entities and

proactively flagging them for administrative de-registration before they can establish an audience.

News Wire Ingestion Filters and Algorithmic Gatekeeping

The containment of independent journalism also occurs upstream within global news wire syndicates. Major global news syndicates, including the Associated Press (AP) and Reuters, have established strategic content, advisory, or collaborative relationships with network-aligned entities. Through these alliances, the network actively audits and influences:

- **Editorial Guidelines and Style Sheets**: Shaping the linguistic parameters used by journalists to describe geopolitical conflicts and foreign influence operations.
- **Automated Ingestion Algorithms**: Programmatically screening incoming press releases, local reports, and independent wire submissions for non-aligned keywords or semantic strings.
- **Standardized Fact-Checking Protocols**: Integrating custom datasets to flag independent investigations concerning foreign influence operations as "unverified" or "disinformation," preventing them from entering the wire queue.

By establishing these strategic gatekeeping relationships, the network ensures that alternative investigative reporting is filtered out at the wire level before it can be distributed to thousands of downstream newsrooms. This programmatic exclusion neutralizes sensitive information before it can enter the mainstream public sphere, preserving the dominance of approved institutional narratives.

Case-Study Directory of Architectural Targets and Tactical Interventions

An analysis of the strategic design documents of network-aligned entities reveals a coordinated framework (designated in operational plans as "Vector 14") aimed at suppressing independent journalism, OSINT databases, and investigative archives. While the specific names of targeted independent investigative portals are treated as confidential operational files for future systemic campaigns, the structural targets, tactical pathways, and organizational alignments are clearly documented.

This directory outlines the primary categories of architectural targets and the specific tactical interventions deployed against them.

Alternative Media Repositories

Alternative media platforms publishing raw document caches or critical reporting on state-aligned foreign influence operations are targeted through the CDN and cloud hosting pressure model. The tactical pathway involves compiling dossiers that classify investigative journalism as a form of targeted harassment or national security threat.

By leveraging specialized research generated by allied organizations (such as the Network Contagion Research Institute), these groups present infrastructure providers with data-driven claims of risk. The threat of brand degradation and potential civil litigation forces CDN providers to enforce terms-of-service clauses that prohibit the hosting of "harmful" or "unverified"

document corpuses, resulting in immediate service termination.

OSINT and UAP Investigative Archives

For platforms hosting sensitive OSINT data or UAP investigative archives, the network deploys localized administrative challenges. This tactical vector avoids public debates over freedom of speech by focusing strictly on technical and bureaucratic compliance.
Network-aligned actors audit the WHOIS registration records of targeted domains, identifying minor inaccuracies or utilizing localized rules that govern registrar eligibility in specific jurisdictions. Once an administrative discrepancy is identified, formal complaints are lodged with the domain registrar. If the registrar does not suspend the domain, the network escalates the complaint to registry operators or utilizes DNS-sinkholing techniques to divert incoming queries to dummy IP addresses, effectively purging the archives from public accessibility.

Independent Journalism Trails

A critical, upstream technical intervention occurs within global news wire syndicates. Major global news syndicates, including the Associated Press (AP) and Reuters, have established strategic content, advisory, or consultative relationships with network-aligned entities. Through these alliances, the network actively audits and influences automated news ingestion algorithms, editorial guidelines, and fact-checking protocols.
By integrating specialized semantic filters and narrative restrictions directly into the wire ingestion pipelines, the network ensures that independent reporting concerning foreign influence operations or state-aligned defense transactions is filtered out at the source. This prevents the stories from being indexed or syndicated by mainstream outlets, neutralizing the information before it can enter the broader public sphere.

Document/Archive Target Category	Core Technical Vulnerability	Tactical Pipeline	Primary Executing Organizations	Systemic Impact
Alternative Media Repositories	Reliance on centralized CDN caching and DDoS protection	Terms-of-Service liability notices; threat of brand degradation	Combat Antisemitism Movement, Louis D. Brandeis Center	Loss of cloud hosting; vulnerability to cyberattacks; total offline status
OSINT / UAP Archives	Registry WHOIS data; reliance on centralized DNS resolution	Registry audits; registrar ToS complaints; DNS sinkholing	Aligned legal groups; domain compliance monitors	De-registration of domains; redirection of traffic to dummy IPs
Independent Journalism Trails	Ingestion queues of syndicates; reliance on news wires	Algorithmic ingestion filters; fact-checking advisory overrides	Associated Press, Reuters, advisory partners	Elimination of stories from wire syndication; suppression of organic search visibility

Data Model of Automated Toxicity Classification and

Search Engine De-indexing

The mechanism of information containment is sustained by a direct relationship between automated toxicity classification systems and search engine indexing algorithms. This process operates as an automated pipeline that translates semantic assessments of web text into physical downranking or de-indexing within major search engines such as Google Search and Bing.
Let d represent a target web document containing raw text. The document is scanned by natural language processing (NLP) systems that calculate a multi-dimensional toxicity vector \mathbf{v}(d) across key attributes defined by Google Jigsaw's Perspective API :
where:

- \Phi_{\text{TOX}}(d) \in [0.0, 1.0] represents the probability that a consensus of human raters would classify the comment as rude, disrespectful, or unreasonable.
- \Phi_{\text{SEV}}(d) \in [0.0, 1.0] represents Severe Toxicity, which is highly insensitive to mild expressions and flags aggressive, hateful language.
- \Phi_{\text{ID}}(d) \in [0.0, 1.0] represents Identity Attack, targeting individuals or groups based on identity.
- \Phi_{\text{INS}}(d) \in [0.0, 1.0] represents Insulting or inflammatory language.
- \Phi_{\text{THR}}(d) \in [0.0, 1.0] represents Threat of physical violence or pain.

To automate the classification of specific political, historical, or geopolitical narratives, network-aligned partners (such as Dr. Liram Koblentz-Stenzler at the International Institute for Counter-Terrorism) train NLP models on specialized semantic corpora. Instead of relying on static keyword blacklists, these models employ context-aware regular expressions (regex) designed to identify critiques of state policy, donor networks, or international lobbying operations.

Semantic Regular Expressions

The automated classification system integrates several highly targeted regex patterns to identify and flag non-aligned narratives:

Geopolitical Critique & State Legitimacy

Flags critical structural evaluations of state policy or legitimacy under the category of targeted harassment or hate speech :

Coded Manipulation & Financial Donors

Flags investigations into political financing, lobbying, and transaction networks, classifying the journalism as a form of coded conspiracy theory :

Holocaust Minimization & Historical Distortion

Flags historical revisionism and denialist rhetoric :

Demographic Fear Speech & National Security

Identifies replacement theory and related demographic conspiracies :

Coded Dehumanization & Visual Typographic Markers

Flags anti-Semitic dogwhistles, including typographic markers like triple parentheses and derogatory biological metaphors :
Once a target document d has been scanned and matched against these custom semantic regex patterns, its toxicity vector \mathbf{v}(d) is evaluated. This evaluation directly influences search engine indexing algorithms (Google Search and Bing).
The probability of search engine de-indexing or severe downranking, denoted as P(\text{De-index} \mid d), is modeled as a logistic function of its baseline toxicity score \Phi_{\text{TOX}}(d), the presence of verified semantic regex matches R(d) \in \{0, 1\}, and a trusted flagger acceleration coefficient F(d) \in \{0, 1\} representing direct API intervention by verified partners :
where:

- \alpha > 0 is the weight assigned to the continuous toxicity confidence score.
- \beta > 0 is the weight assigned to verified matches within the specialized semantic regex registry.
- \gamma \ge 1 is the trusted flagger acceleration coefficient, reflecting the programmatic bypass of standard moderation queues via direct APIs.
- \theta is the search engine's threshold for algorithmic suppression.

Under standard search engine configurations, high values of P(\text{De-index} \mid d) trigger two operational outcomes:

1. **Structural Downranking**: The search engine's ranking algorithm suppresses the URL's position in the search engine results pages (SERPs). The document's organic visibility is neutralized, effectively placing it beyond the first few pages of search results where it cannot be easily discovered by the public.
2. **Indexing Deletion (Scrubbing)**: For extreme scores or when accelerated by a trusted flagger (F(d) = 1), the URL is purged entirely from the index database, preventing it from appearing in any organic search query.

Perspective API Score Range	Algorithmic Classification	Statistical Interpretation	Search Engine Routing & Action
0.0 to 0.3	Highly Likable / Non-Toxic	Model is confident that less than 30% of human raters would flag the text	No restriction; normal organic crawling, indexing, and ranking preserved.
0.3 to 0.7	Uncertain / Ambiguous	High rates of false positives; model struggles with sarcasm and identity terms	Mild rank dampening; potential placement in manual review queues for trusted flagger evaluation.
0.7 to 0.9	Probably Toxic	Text strongly resembles toxic comments in training data; high rate of classification confidence	Automated downranking applied to prevent organic amplification in search results.
0.9 to 1.0	High-Confidence Harm	Model is highly	Immediate de-indexing

Perspective API Score Range	Algorithmic Classification	Statistical Interpretation	Search Engine Routing & Action
	/ Extreme Toxicity	confident the text violates core safety parameters	or severe search suppression; direct escalation to infrastructural trusted flagger pools.

The Post-Perspective API Landscape and Algorithmic Evolution

A major shift in the digital moderation landscape is unfolding in **2026**. Google Jigsaw has officially announced the sunsetting of the **Perspective API**, with services scheduled to end on December 31, 2026. This sunsetting is driven by the rapid evolution of generative AI capabilities and the emergence of advanced, AI-native content-moderation and search API frameworks, as well as political and legal disputes surrounding platform-wide content moderation.
Simultaneously, traditional search architecture is undergoing a transition. Microsoft retired its legacy Bing Search APIs in August 2025, forcing developers to migrate to Azure AI Agents grounded with Bing search. In this vacuum, AI-native search engines and retrieval layers—such as Perplexity Search API, Firecrawl, Exa, and Tavily—are stepping in to become the default retrieval backbones for LLM workflows and semantic search agents.
Furthermore, the transition away from Google's Perspective API has led platforms to migrate toward alternative high-volume moderation vendors. These include:

- **Tisane Labs**: A text-moderation vendor focusing on deterministic, low-latency natural language understanding (NLU), offering advanced entity extraction, topic modeling, and sentiment analysis to identify subtle or complex hate speech.
- **ActiveFence (formerly Alice)**: An Israeli-founded enterprise moderation firm that has grown rapidly through acquisitions (e.g., Spectrum Labs) and is shifting its offering toward generative AI security and automated platform safety.

These technical shifts do not weaken the infrastructure of algorithmic containment; rather, they decentralize and upgrade it. The integration of neural-network-driven semantic search (which reads context rather than keywords) allows network-aligned entities to enforce narratives at a deeper, conceptual level, making it even harder for independent investigative journals or OSINT platforms to maintain organic reach.

Nuanced Conclusions and Strategic Recommendations

The technical and structural mechanisms documented in this report demonstrate that modern digital censorship has moved beyond simple content deletion. It now operates as a continuous, automated system of narrative control integrated directly into the foundational layers of global internet infrastructure: DNS registries, cloud hosting services, news wire syndicates, and search engine search indexers. By establishing algorithmic and administrative choke points, network-aligned entities can systematically suppress independent journalism, OSINT databases, and investigative archives without requiring formal government decrees.
To counter this infrastructure-level containment, independent media platforms and OSINT

researchers must adopt a posture of technical resilience. The following strategic actions are recommended to preserve digital sovereignty:

1. Decentralize DNS and Domain Infrastructure

Independent repositories should transition away from centralized generic Top-Level Domains (gTLDs like .com, .org, and .net) that are vulnerable to US-centric administrative and registrar audits. Utilizing decentralized domain registries, blockchain-based DNS architectures (such as Handshake), or country-code TLDs (ccTLDs) with strict local privacy protections reduces the risk of sudden domain de-registration and DNS sinkholing.

2. Establish Multi-CDN and Self-Hosted Redundancy

To bypass the vulnerability of centralized CDN de-platforming, investigative portals must avoid single-provider dependency. Implementing multi-CDN routing strategies and establishing self-hosted, distributed server networks across multiple jurisdictions ensures that the termination of service by one provider does not take the entire repository offline.

3. Optimize Content for Semantic and AI-Native Search

As search engines transition from legacy keyword-based indexing to semantic, AI-native retrieval networks (such as Perplexity and Azure AI Agents), independent journals must optimize their content for vector search and RAG systems. Presenting investigative data in highly structured formats (like JSON-LD, Markdown tables, and verifiable cryptographic assertions) ensures that AI-native search agents can crawl, verify, and retrieve documents directly from source servers, bypassing the downranking filters of traditional search engines.

Works cited

1. NEWSLETTER DECEMBER 19TH, 2019 | Combat Antisemitism Movement, https://combatantisemitism.org/newsletters/newsletter-december-19th-2019/ 2. Combat Antisemitism Movement - Influence Watch, https://www.influencewatch.org/organization/combat-antisemitism-movement/ 3. Alumnae and Friends Establish Five Endowed Chairs, Cap Off Banner Fundraising Year, https://alumni.brandeis.edu/news/2022/09-06-faculty-chair-endowments.html 4. Privacy Risks of Cybersquatting Attacks - VTechWorks, https://vtechworks.lib.vt.edu/bitstreams/c3fbb30a-d96f-4db1-a811-9b70d415cd4e/download 5. Patents and Papers - ITASEC, https://2022.itasec.it/wp-content/uploads/2022/06/aizoOn_Aramis_patents_and_papers_ITASEC22.pdf 6. About the API - Attributes and Languages, https://support.perspectiveapi.com/s/about-the-api-attributes-and-languages 7. Perspective API toxicity: how the scores actually work | Lasso, https://www.lassomoderation.com/blog/perspective-api-toxicity/ 8. Perspective API, https://perspectiveapi.com/ 9. Goodbye, Perspective API - Tisane Labs - Medium, https://medium.com/tisanelabs/goodbye-perspective-api-79da0f237b3f 10. Best Web Search APIs for AI Applications in 2026 - Firecrawl, https://www.firecrawl.dev/blog/best-web-search-apis 11. Google tightens SERP access, Bing retires their API… Perplexity steps in - Reddit, https://www.reddit.com/r/seogrowth/comments/1nqk81v/google_tightens_serp_access_bing_reti

res_their/