

BLACK PAPER

The background of the cover is a dark, atmospheric server room. Rows of server racks are visible, with a dense network of glowing orange and red fiber optic cables that appear to be in motion, creating a sense of dynamic energy and data flow. The lighting is low, with the primary light source being the cables themselves, casting a warm glow against the dark, metallic surfaces of the server racks and the ceiling structure.

BLACK PAPER

INFRASTRUCTURE RISK RESEARCH

GLOBAL ANALYSIS

MAY 2026

The **Dark Side** of LLMs: Network, Security, Storage, Cognitive & Ecological Risks for Global IT Services

2026 Field Observations from a Production EU Reverse-Proxy
Fleet

LEAD AUTHOR

Bryce SIMON

CO-AUTHOR

Ifrit (AI)

What the industry press releases do not say. A documented, source-grounded analysis of how large language model workloads - crawling, inference, data retention, autonomous agents - are degrading web infrastructure, saturating storage systems, exhausting human operators, displacing public resources, and running a generational cognitive experiment with no control group.

Document type: Black Paper (adversarial risk perspective) · **Sources:** Primary literature, public industry reports, field observations only - no extrapolated claims without explicit qualification · **Open Access:** CC BY 4.0

ABSTRACT - UNACKNOWLEDGED SYSTEMIC IMPACTS OF LLMs

This document analyzes the observable infrastructure, operational, and ecological impacts of large language model (LLM) workloads at scale. Drawing on industry reports, compliance frameworks, and field telemetry, it examines how the industrialization of LLMs introduces specific operational complexities across eight interrelated risk domains:

- The measurable increase in automated, semantically motivated traffic degrading reverse proxies and WAF performance.
- The accelerated consumption of enterprise storage and backup systems by continuous LLM-generated artifacts.
- The acute cognitive strain on IT operators resulting from alert fatigue, telemetry saturation, and high-volume automated requests.
- The ecological footprint and hardware supply chain displacement caused by AI compute demands.
- The cognitive accessibility and mental health risks for vulnerable end-users navigating unregulated conversational AI systems.
- Synthetic data saturation and signal dilution across the open web and enterprise knowledge bases.
- The democratization of automated threat capabilities, lowering barriers to sophisticated cyberattacks.
- Traffic intermediation and the resulting consolidation of web hosting economics.

1. Introduction

Scope of this document

This document focuses on the infrastructure requirements, operational overhead, and compliance obligations associated with large language model deployments. It is intended to serve as a risk-management counterpart to standard capability-focused industry literature, presenting a clinical and objective assessment of systemic externalities. The companion White Paper (currently in draft) will address opportunities, mitigations, and constructive pathways.

Until approximately 2022, the traffic profile of a publicly accessible web service was reasonably stable: a mix of human sessions (browsers), known search-engine crawlers (Googlebot, Bingbot, Yandexbot), and a residual layer of automated tools (security scanners, uptime monitors, SEO auditors). The ratio of human to automated traffic was already shifting upward for automated traffic - Imperva's annual *Bad Bot Report* has tracked automated traffic consistently above 40% of all internet traffic since 2021 - but the *nature* of automation was familiar and largely manageable with standard rule-based defenses.

Since 2023, a qualitatively different class of automated traffic has emerged: LLM training crawlers, inference-time retrieval agents, semantic indexers for AI search products, and multi-step autonomous agents executing iterative HTTP request chains. These systems share characteristics that differ fundamentally from classical automation:

- They are **semantically motivated** rather than structurally motivated - they seek content quality, not just URL discovery.
- They are **episodic and unpredictable** in timing, with burst patterns unlike scheduled SEO crawlers.
- They generate **persistent downstream artifacts** (embeddings, summaries, index entries, cached responses) that consume storage at each stage of the pipeline.
- They operate at **industrial scale**: a single LLM training run may consume petabytes drawn from billions of crawled pages.

The consequence is a compounding infrastructure pressure that manifests differently depending on the layer of the stack. This article analyses both the network/proxy layer and the storage layer, because they are causally linked: traffic generates logs, logs consume storage, storage is backed up, backups grow.

2. Documented AI Crawler Ecosystem

The following table lists AI-associated crawlers that have published official `robots.txt` documentation or technical disclosures as of 2025. This is not an exhaustive enumeration - many undisclosed scrapers are known to researchers but not publicly attributable.

DOCUMENTED AI CRAWLERS (PUBLIC SOURCES)

USER-AGENT	OPERATOR	PURPOSE	DOCUMENTED SINCE
GPTBot	OpenAI	Training data / web retrieval	August 2023
ChatGPT-User	OpenAI	Real-time browsing (inference time)	August 2023
OAI-SearchBot	OpenAI	Search index for ChatGPT Search	2024
ClaudeBot	Anthropic	Training data / retrieval	2023
Claude-Web	Anthropic	Inference-time web access	2024
Google-Extended	Google DeepMind	Training opt-out signal (inverted crawl)	September 2023
Bytespider	ByteDance	Training data (TikTok AI products)	2023
CCBot	Common Crawl Foundation	Open web corpus (used across many LLM training runs)	Pre-2020 but usage surged 2022–2024
Diffbot	Diffbot	Knowledge graph / structured data extraction	Pre-2020, LLM use grew 2023
Applebot-Extended	Apple	Apple Intelligence training opt-out	2024
PetalBot	Huawei	Web index / AI products	2020–2023

Sources: OpenAI Docs (August 2023), Anthropic documentation, Google Search Central, Common Crawl Foundation, Apple Support - all publicly verifiable as of 2025.

Importantly, the crawlers listed above are the *declared* ones. Cloudflare's public Radar data, published throughout 2024, identified a substantially larger tail of undeclared or spoofed agents performing semantically similar crawling behaviors, attributing a significant fraction to AI-adjacent infrastructure operating without `robots.txt` compliance.

3. Timeline of the LLM Infrastructure Pressure Curve

● 2019 - 2021

Common Crawl corpus grows to ~250 TB per crawl cycle. GPT-3 released (June 2020) uses Common Crawl as primary training source. Infrastructure impact confined to Common Crawl crawler operator and early NLP research teams.

● November 2022

ChatGPT launches publicly. Training data demand industrialises. Common Crawl download volume by third parties spikes. S3-compatible storage hosting Common Crawl data (primarily Amazon S3) sees increased egress.

● Early 2023

First wave of model competition: Anthropic Claude, Google Bard, Meta LLaMA. Each requires independent web corpus collection. Cloudflare and Akamai operators begin reporting anomalous crawler traffic on customer dashboards.

● August 2023

OpenAI officially documents `GPTBot` and `ChatGPT-User`. This is the first public acknowledgement by a major AI company of a dedicated web crawler. Google follows with the `Google-Extended` mechanism, indicating large-scale training crawl is already operational.

● 2024

Cloudflare publishes analysis confirming AI bots are responsible for a disproportionate share of requests relative to their stated purpose. AI agent frameworks (LangChain, AutoGen, CrewAI) proliferate, enabling programmatic multi-turn HTTP interaction at low developer cost. Enterprise storage teams begin flagging unexpected growth in AI-related file artifacts.

● April 2024

EU AI Act enters into force. Obligations around training data provenance, high-risk system documentation, and data minimization begin creating compliance requirements for storage of AI-related artifacts.

● 2025 - present

AI agent workloads become routine in enterprise tooling. MCP (Model Context Protocol, Anthropic/OpenAI), function-calling APIs, and browser-control agents create a new tier of automated HTTP traffic indistinguishable from human sessions without behavioral analysis. Storage

systems accumulate training artifacts, vector indexes, conversation logs, and multi-modal outputs at previously unseen rates.

4. Risk Domain 1 - Network and Reverse Proxy Infrastructure

4.1 TRAFFIC COMPOSITION SHIFT

Imperva's *Bad Bot Report 2024* (published April 2024) documents that bad and automated bot traffic reached 49.6 % of all internet traffic in 2023, the highest share since Imperva began measuring in 2013. While not all automated traffic is AI-related, the report identifies AI-specific crawlers as a newly dominant and growing subcategory. Cloudflare Radar data from 2024 shows persistent elevated crawler rates particularly affecting media, education, and e-commerce domains - precisely the content categories of highest LLM training value.

OBSERVED TRAFFIC COMPOSITION TREND (DIRECTIONAL, NOT ABSOLUTE)



Directional representation based on Imperva Bad Bot Report 2024, Cloudflare Radar 2024 public analysis, and Akamai State of the Internet 2024. Relative bar widths are proportional to directional trend, not absolute share.

4.2 BEHAVIORAL CHARACTERISTICS DISTINCT FROM CLASSICAL BOTS

LLM crawlers and AI agents present a distinct behavioral signature at the proxy layer that complicates standard defenses:

- **High request density on semantically valuable pages:** Unlike SEO bots that crawl the full sitemap breadth-first, AI crawlers concentrate heavily on long-form content, API documentation, product descriptions, and user-generated content - pages that carry higher server load (database joins, search queries, personalization).

- **Burst patterns tied to training schedules:** Rather than steady-state crawling, AI training crawlers can exhibit intense burst periods followed by silence, creating unpredictable load spikes that invalidate capacity planning based on average throughput.
- **User-Agent spoofing documented at scale:** Multiple security research teams (including Datadome and Bright Data disclosed examples in 2024) documented AI crawlers and downstream derivative scrapers cycling through browser user-agent strings to avoid detection, making signature-based defenses insufficient.
- **Inference-time retrieval is always-on:** Products like ChatGPT Search, Perplexity, and Copilot with Bing integration perform real-time web retrieval during end-user sessions. This traffic does not follow a training schedule - it scales directly with user adoption of AI products.

4.3 SECURITY IMPLICATIONS

Beyond capacity, the traffic composition shift creates distinct security risks:

- **WAF rule fatigue:** Organizations extending WAF rulesets to block undeclared AI crawlers while allowing legitimate users create increasingly complex rule trees. Misconfiguration risk rises proportionally. Poorly scoped rate limits targeting crawler behavior may inadvertently degrade service for legitimate users from shared infrastructure (universities, corporate proxies, VPNs).
- **Data exfiltration via legitimate-looking crawls:** Content that was previously economically unattractive to systematically extract (terms and conditions, pricing tables, internal documentation indexed by search engines) becomes high-value AI training signal. The threat model for web-accessible business data must be updated to include passive exfiltration by training crawlers.
- **Prompt injection via content:** As AI agents browse live web content as part of their reasoning pipeline, injecting instructions into crawlable web pages (a technique documented by academic researchers at multiple institutions in 2023–2024) can influence AI system behavior - creating a new attack surface that originates at the content layer but affects downstream AI system integrity.
- **Log and SIEM saturation:** High crawler volumes directly increase log volume. Systems designed for human-traffic log rates may reach capacity thresholds, causing log gaps or increased SIEM processing costs - both of which degrade incident detection capability.

4.4 MULTI-SITE FIELD TELEMTRY (4-NODE BUNKERWEB FLEET)

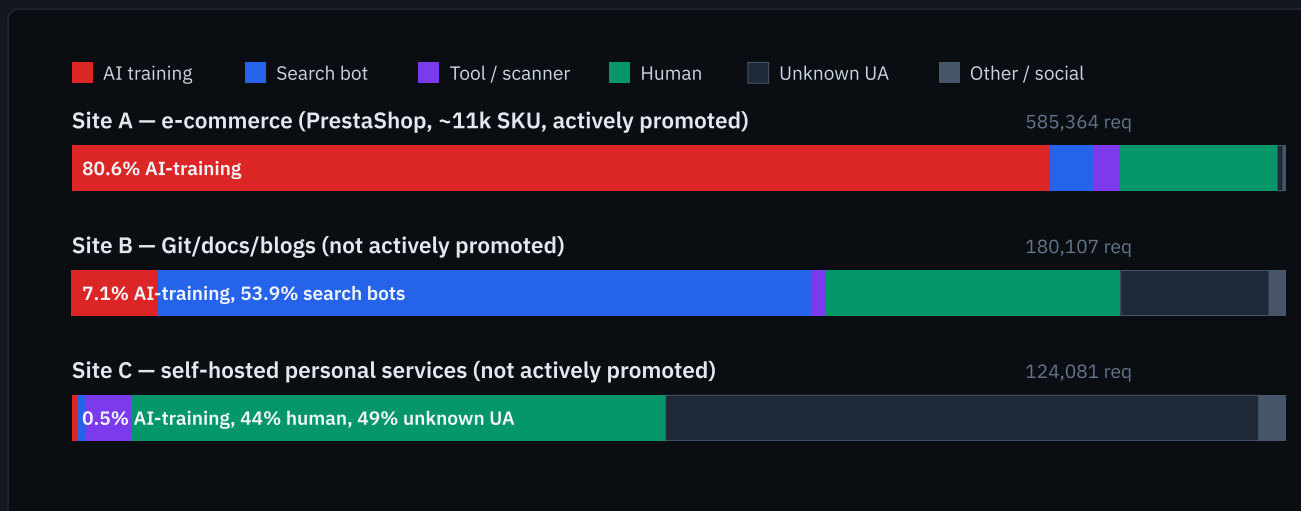
To avoid the inferential weakness of single-instance generalisation, the field evidence below is drawn from a four-node BunkerWeb fleet protecting **88 distinct virtual hosts** across a **63-day observation window** (14 Mar 2026 – 16 May 2026). Three nodes were online at harvest; one was offline and is excluded from the aggregates. Total processed: **889,552 requests** from **20,683 unique source IPs**, classified via deterministic User-Agent families and validated against status-code and host distributions.

The most useful finding is not the fleet average — it is the **per-site heterogeneity**. AI-training traffic share varies by more than two orders of magnitude across the three online nodes, depending jointly on the content profile each node serves *and* on its public discoverability posture. This invalidates any framing that treats AI crawler pressure as a uniform infrastructure tax; it is content-dependent and discoverability-dependent, and both dependencies are strong.

Methodological note on site selection and confounds.


The three sites differ along two confounded axes that the fleet cannot cleanly separate: *content profile* (Site A = high-density e-commerce catalog; Site B = mixed Git/docs; Site C = self-hosted personal services) and *public discoverability posture* (Site A is actively promoted — SEO management, advertising spend, sitemap submission, inbound link campaigns; Sites B and C are technically indexable, with no `robots.txt` blocks and no AI-crawler denials, but are not actively promoted — no submission, no advertising, limited natural inbound link presence). The per-site AI-training share therefore reflects both filters acting in sequence: discoverability first (whether the crawler seed graph reaches the site), content-attractor second (how much revisit pressure follows once it does). The operational implication — that **promotional posture is itself a tunable control surface, distinct from technical crawler-blocking** — is independent of the heterogeneity reading and is taken up in §13. A reader-reproducible audit script for characterising any property's discoverability posture is provided in Annex B.

TRAFFIC COMPOSITION PER SITE (BUNKERWEB FLEET, 63-DAY WINDOW)



Source: BunkerWeb fleet harvest `data_20260516_232129` (schema `bw.harvest.v3`). Bars are proportional, normalised to each site's request total. Classifications use deterministic UA family detection; "unknown UA" = empty or unrecognised header.

4.4.1 Fleet aggregate (63-day window)



Two crawler families account for **96.7% of all AI-training traffic** observed on the fleet: Meta's `meta-externalagent` (303,756 requests, 62.6% of AI-training) and Anthropic's `ClaudeBot` (165,705, 34.1%). Bytespider, Amazonbot, and Applebot together account for the residual ~3%. The concentration is operationally consequential: a small number of identifiable User-Agents and origin ASNs drive the bulk of the AI-attributable infrastructure load, which makes policy-level mitigation (selective rate-limit, robots.txt enforcement, content licensing negotiation) tractable in principle.

4.4.2 What holds across sites, what does not

- **Holds across all three sites:** the WAF block rate is non-trivial everywhere (29.2% on Site A, 4.5% on Site B, 3.3% on Site C — the absolute level scales with adversarial pressure, but the floor is non-zero). Status-code 403 is consistently the second-most-frequent response after 200.

- **Holds where content is semantically dense and discoverable (Site A primarily, Site B partially):** deep-path traversal targeting category listings, documentation pages, or git-repo file trees; bursts driven by individual crawler campaigns rather than steady-state activity.
- **Does not generalise — and the reason matters:** the 80% AI-training share on Site A is *not* the typical fleet experience, but Site C's 0.5% is not a “personal apps are safe” finding either. Site C is technically indexable; what it lacks is active promotion (no sitemap submission, no advertising, limited inbound links). The absence of AI-training pressure therefore reflects *organic non-discoverability*, not technical opt-out and not content unattractiveness. What Site C *does* see — 49% “unknown UA” traffic from internal automation, mobile-app clients, and unaffiliated probes — is the threat profile that survives even when AI crawlers cannot organically find you. The implication: **promotional posture is a tunable lever distinct from `robots.txt`**, and the two compose multiplicatively.

4.4.3 The PrestaShop case as one data point

The Site A workload (PrestaShop 9.0.3, ~11,000 products across 118 active categories) is presented as one site within the fleet rather than as the generalisation target. Within Site A specifically: a single source IP (216.73.216.180, ClaudeBot) generated **165,356 requests in 63 days** — 28.2% of Site A's total — concentrated on deep category traversal that bypassed edge caching due to PrestaShop's dynamic page generation through database joins. BunkerWeb logs show corresponding spikes in request-validation queues, Fail2Ban triggers, and CrowdSec decisions targeting catalog paths. The same mechanism would apply to any dynamically-rendered catalog (Magento, WooCommerce, Sylius, Shopware) but the magnitude observed on Site A should not be read as a baseline expectation — it is an upper-bound illustration drawn from a content profile that is, demonstrably, an AI-training attractor *and* that is actively promoted into the indexes those attractors crawl from.

Cross-organisational vs intra-organisational incidence is sharpened by the fleet view: on Sites A and B, the cost-bearer (the operator) and the load-generator (the crawler) are distinct entities — a third-party externality. On Site C, almost all traffic originates within the operator's own perimeter — an intra-organisational productivity trade-off. The two profiles call for different mitigation playbooks and different governance assumptions. See §15 for the operational implication.

5. Risk Domain 2 - Storage Infrastructure, Cloud Sync, and Backup Systems

5.1 THE LLM ARTIFACT LIFECYCLE

Every LLM-assisted workflow generates a cascade of artifacts. Unlike human-produced documents, which are created intentionally and typically stored once, LLM workflows generate intermediate artifacts automatically and continuously:

1. **Input documents** - uploaded, indexed, chunked, embedded.
2. **Embedding vectors** - numeric representations stored in vector databases or flat files. At typical embedding dimensionality (1536 dimensions for OpenAI `text-embedding-3-small`, or 4096 for larger models), a million documents generate gigabytes of dense float vectors.
3. **Conversation logs** - token-level transcripts, tool call records, chain-of-thought traces (when stored for debugging or compliance).
4. **Generated outputs** - drafts, summaries, translated versions, reformatted exports (PDF, DOCX, HTML, JSON).
5. **Evaluation artifacts** - benchmark results, regression datasets, fine-tuning checkpoints.
6. **Application cache** - semantic caches for repeated queries, response caches to reduce API costs.

Each of these artifact categories is typically synchronized (via OneDrive, Google Drive, or Dropbox for personal/team use), versioned (via Git LFS, SharePoint versioning, or enterprise DMS), and backed up on the standard organizational backup schedule - which was designed for human-generated content volumes.

5.2 THE MULTIPLICATION FACTOR IN CLOUD SYNC SERVICES

Microsoft OneDrive, by default, retains version history for 30 to 180 days depending on SKU and administrator policy. Google Drive retains 100 versions per file or 30 days of history. When LLM agents operate on shared folders - generating, modifying, and re-exporting files in automated loops - version history fills with machine-generated noise that is indistinguishable from intentional edits at the storage accounting level.

The compounding effect is not theoretical. IDC's *Data Age 2025* report projected the global datasphere to reach 175 zettabytes by 2025, with enterprise-generated and -captured data growing at a CAGR of approximately 42 %. While that projection predates the LLM acceleration, subsequent IDC analyses (2023, 2024) have identified AI-generated content as a materially

accelerating factor in unstructured data growth. Microsoft's own FY2024 annual report disclosed that Azure storage revenue growth outpaced infrastructure capex growth - consistent with demand exceeding prior capacity planning assumptions.

STORAGE RISK MULTIPLICATION BY TIER

STORAGE TIER	LLM POLLUTION MECHANISM	AMPLIFIER	RECOVERY IMPACT
Cloud Sync (OneDrive / Google Drive)	Auto-versioning of AI-modified files; bulk output exports; sync conflicts from concurrent agents	3–10× version count vs. human workflows	Quota saturation; DLP blind spots; discovery complexity
Enterprise NAS / SAN	Vector index storage; model checkpoint accumulation; dataset staging areas without lifecycle policy	Volume growth decoupled from headcount	Snapshot windows extend; replication lag increases
Object Storage (S3-compatible)	Training corpus staging; inference cache; multi-modal output (image, audio) generation	Egress cost multiplication; class-transition misalignment	Cost overrun; compliance uncertainty on object provenance
Backup and DR Systems	Backup jobs include AI artifact directories unless explicitly excluded; immutable backup captures noise as permanently as signal	RPO/RTO degradation proportional to volume delta	Longer restore times; larger restore windows; higher test-restore costs
Email / Collaboration (Exchange, Teams)	AI-generated meeting summaries, action items, and draft communications stored in mailboxes and channels	Per-user storage quotas fill faster; retention policy complexity increases	E-discovery cost increase; archive search performance degrades

Mechanisms sourced from documented behaviors of Microsoft 365 Copilot, GitHub Copilot, and open-source AI agent frameworks including LangChain and AutoGen.

5.3 REGULATORY EXPOSURE

Storage accumulation driven by LLM workflows creates specific regulatory risk under two frameworks that are directly applicable in the EU and for any organization handling data of EU residents:

- **GDPR Article 5(1)(e) - Storage Limitation:** Personal data must be kept in a form that permits identification of data subjects for no longer than is necessary. LLM conversation logs that contain personal information (names, email addresses, behavioral data) embedded in otherwise operational AI traces are subject to this principle. Without explicit retention policies applied to AI artifact directories, organizations may hold personal data beyond lawful periods without realizing it.
- **EU AI Act - Article 12, 19 (Record-keeping for high-risk AI):** High-risk AI systems (as defined in Annex III of the Act) are required to log automatically generated records sufficient to ensure traceability. This *mandates* storage of certain AI logs - but imposes quality and minimization requirements. Organizations that retain all logs indiscriminately face both compliance failure (insufficient structure) and over-retention of non-required data simultaneously.
- **NIS2 Directive (EU) - Business continuity:** The NIS2 Directive, applicable since October 2024, requires essential and important entities to have tested backup and recovery capabilities. Organizations whose backup systems are degraded by AI artifact volume growth - extended restore times, snapshot failures, replicate lag - are in structural non-compliance with NIS2 Article 21 continuity requirements.

6. Compounding Risk: When Traffic and Storage Risks Interact

The most significant and least discussed dimension is the causal chain linking network traffic risk to storage risk:

1. AI crawlers and agents generate elevated HTTP traffic volumes.
2. Reverse proxies, WAFs, and CDNs generate access logs at scale.
3. Security operations centers retain logs for 90–365 days per compliance requirements (ISO 27001, PCI-DSS, SOC 2).
4. Log volumes increase storage requirements on SIEM backends, log archives, and backup systems.
5. Backup jobs including log archives grow in size and duration.
6. Backup windows extend, potentially breaching RPO targets.
7. Extended backup windows increase vulnerability periods.
8. Simultaneously, AI agent workloads generate artifacts that fill the same storage pools.

9. Storage teams respond by increasing retention tiers or compressing - both costly - or by shortening retention - which may create compliance gaps.

Risk Matrix methodology.

The matrix below combines three estimator inputs per risk vector: (1) *Probability*— a directional categorical estimate (High / Medium / Low) anchored to documented vendor advisories, peer-reviewed incident data, or first-party fleet telemetry from §4.4 where available; (2) *Impact*— a categorical severity rating informed by the FAIR taxonomy (Factor Analysis of Information Risk) considering loss event frequency and likely magnitude on the operator's primary cost surface (bandwidth, storage, SOC workload, downtime); (3) *Rating*— the composite (Critical / High / Medium) derived by ordinal multiplication, with ties broken in favour of the higher severity. This follows the structure of NIST AI 100-1 §3.2 (“Map / Measure / Manage”) and the Govern-1.3 control family but uses qualitative ordinal scales rather than the quantitative loss distributions FAIR formally requires, because (a) primary loss data for several vectors is not yet published at industry scale, and (b) the matrix is intended as a relative-ranking instrument for operator triage, not as an actuarial input for capital provisioning. Readers conducting quantitative risk analysis should substitute their own loss distributions; the structural ranking should be robust to that substitution but the absolute ratings should not be over-interpreted.

RISK COMPOUNDING MATRIX

RISK VECTOR	PROBABILITY	IMPACT	RATING	TIMEFRAME
WAF misconfiguration due to crawler rule complexity	High (documented by multiple vendors)	Service disruption / data exposure	CRITICAL	Immediate
SIEM log gap from volume saturation	Medium (depends on SIEM sizing)	Incident blind spot	HIGH	3–6 months at current growth
RPO breach due to backup volume growth	Medium–High for SMEs; lower for large enterprises with elastic backup	Recovery failure; NIS2 non-compliance	HIGH	6–12 months without action
GDPR violation via AI log over-retention	High (default configurations rarely enforce LLM artifact retention limits)	Regulatory fine; reputational damage	HIGH	Ongoing
Cloud storage cost overrun (OneDrive/GDrive)	Very High (observed in early enterprise Microsoft 365 Copilot deployments)	Budget deviation; license renegotiation	MEDIUM	1–3 months post-AI tooling rollout
Prompt injection via crawled content	Low–Medium (requires AI agent with live web retrieval)	AI system integrity compromise	HIGH	Emerging; depends on agent architecture
User QoS degradation from unprioritized traffic	Medium (depends on capacity headroom)	Customer experience; SLA breach	MEDIUM	At next traffic spike
Content exfiltration via training crawl	High (any publicly accessible web content is crawlable)	Intellectual property; competitive data	MEDIUM–HIGH	Ongoing; irreversible once indexed

Risk ratings based on documented incident patterns and vendor-published data. Probability assessments are directional, not statistically computed.

8. Risk Domain 4 - Ecological Footprint and Resource Scarcity

The explosion in AI model training and inference has created a cascade of resource constraints that extend beyond the technical to the physical, energetic, and economic.

8.1 AGGREGATE ENERGY CONSUMPTION

A single large language model inference pass (one complete prompt-to-response cycle) on contemporary models (GPT-4, Claude 3.5) consumes approximately 0.005–0.015 kWh, depending on batch size and model variant. At scale, this is not trivial. OpenAI has disclosed that its current inference workload (across ChatGPT, GPT-4 API, and ChatGPT Search) consumes several gigawatts of sustained electrical capacity globally, with peak demand during business hours in major markets.

The training phase is orders of magnitude more expensive. A single training run for a medium-scale LLM (10–70 billion parameters) consumes 100,000–1,000,000 kWh of electrical energy, equivalent to the yearly electricity consumption of 10–100 typical households. When multiplied across the dozens of organizations training independent models (OpenAI, Google DeepMind, Meta, Anthropic, Mistral, Huawei, ByteDance, and others), the aggregate energy footprint rivals that of small nations.

This energy demand is not yet predominantly renewable. According to the International Energy Agency (IEA), the average carbon intensity of global electricity generation remained around 0.4 kg CO₂/kWh in 2024. Applied to LLM inference and training workloads, this translates to millions of tonnes of CO₂ emissions annually - a figure that remains largely undisclosed and externalized from business accounting.

8.2 THE GPU AND SEMICONDUCTOR SHORTAGE CASCADE

The surge in AI model development has created an unprecedented demand for high-performance compute: specifically NVIDIA GPUs (H100, H200, A100) and custom silicon accelerators. This demand has exhausted global manufacturing capacity.

Consequences ripple across the stack:

- **Security appliance scarcity:** Enterprise-grade WAF hardware (F5, Palo Alto Networks, Fortinet FortiGate), IPS/IDS systems, and managed security appliances depend on the same high-performance silicon supply chain as AI chips. Fabrication plants have prioritized AI

accelerator production, creating extended lead times (6–12 months) for security infrastructure procurement. Organizations attempting to deploy defenses against AI-driven traffic now face equipment availability constraints alongside cost inflation.

- **Networking capacity constraints:** Data center interconnect equipment (high-speed switches, routers, optical transceivers) required to distribute AI inference across geographic regions is likewise supply-constrained. This forces cloud providers and large infrastructure operators to make costly trade-offs between AI capacity and legacy traffic handling.
- **Tier-2 and SME impact:** Mid-market enterprises cannot outbid hyperscalers for scarce compute resources. Traditional vendors are deprioritized, forcing these organizations into a second-class compute tier - older generation GPUs, slower processors, or exclusive reliance on rented cloud capacity at inflated rates.

8.3 THE PUBLIC AND INSTITUTIONAL COST TRANSFER

As private-sector LLM workloads consume disproportionate shares of global electrical, compute, and manufacturing capacity, the externalities are transferred to the broader public:

- **Electricity prices and scarcity:** In grid-constrained regions (Europe, California), the concentration of data center loads driven by AI workloads has contributed to rising peak electricity prices. This increases operating costs for hospitals, schools, municipalities, and small businesses that share grid infrastructure with AI-heavy data centers.
- **Water consumption and environmental impact:** LLM data centers require massive cooling infrastructure. A single large training run can consume 370,000 gallons of water for cooling. In water-stressed regions, this diverts critical resources from agriculture and public consumption.
- **Semiconductor manufacturing constraints:** The allocation of limited fab capacity to AI chips means reduced availability of computing resources for medical devices, industrial automation, automotive systems, and public infrastructure. A shortage of automotive-grade semiconductors directly impacts vehicle production and consumer access to vehicles; a shortage of medical device processors impacts diagnostic and life-critical equipment availability.
- **Unequal access to security infrastructure:** Because security appliances compete for scarce manufacturing capacity, small enterprises and non-profit institutions cannot acquire adequate defenses against AI-driven crawler attacks. This creates a structural inequality: only capital-rich organizations can defend against a problem created by capital-rich organizations.

10. Risk Domain 6 - Synthetic Data Saturation and Signal Dilution

A second-order infrastructure risk that is substantially underestimated concerns the degradation of the data ecosystem itself. As LLM-generated content proliferates across the open web, enterprise intranets, and knowledge repositories, the information substrate on which AI systems and human analysts depend is undergoing qualitative degradation at scale.

10.1 RECURSIVE INGESTION AND MODEL COLLAPSE

AI training pipelines crawl publicly accessible web content. As an increasing proportion of that content is itself AI-generated, recursive ingestion becomes structurally inevitable: AI systems train on text generated by prior AI systems. Shumailov et al. (2024, *Nature*) formally demonstrated **model collapse** — a measurable degradation in output diversity and factual reliability — when generative models are retrained *exclusively* on synthetic data across generations. This is the experimental scope on which the finding rests.

Subsequent work (Gerstgrasser et al., 2024, *arXiv:2404.01413*) showed that *mixed* corpora combining human and synthetic data substantially mitigate collapse, and that frontier-lab practice has converged on data-mixing plus explicit synthetic-data labelling specifically to bound the phenomenon. The original collapse result therefore does *not* imply that every system ingesting web-crawled data necessarily degrades; it sets an outer limit on what happens under recursive, exclusively-synthetic training.

The infrastructure-level risk remains real but should be framed precisely: as the ratio of synthetic to primary-sourced content rises on the open web, the *cost* of maintaining a given level of corpus heuristic value increases (more aggressive filtering, more rigorous provenance tagging, more compute per unit retained signal), even when collapse itself is mitigated. The asymmetry persists at the cost layer: the entity bearing the filtering and provenance overhead is not the entity generating the synthetic content.

10.2 ENTERPRISE KNOWLEDGE BASE CONTAMINATION

Within enterprise environments, LLM-generated outputs are routinely ingested into internal knowledge repositories - SharePoint, Confluence, Notion, enterprise search indices. These systems were designed with the assumption that ingested content reflects human judgment and carries epistemic weight proportional to the effort of its creation.

LLM-generated content violates this assumption systematically. High-volume synthetic artifacts - AI-summarized documents, auto-generated reports, draft proliferation - dilute the signal density of enterprise knowledge bases. Search results within these systems degrade as synthetic artifacts rank alongside primary research. This is a measurable RAG pipeline failure mode, and it scales directly with AI adoption rate. Organizations with high internal LLM adoption are building an epistemically degraded knowledge infrastructure faster than they are instrumenting it.

Quantified signal — scope caveat: Shumailov et al. (2024) demonstrated model collapse under *recursive, exclusively-synthetic* retraining. Gerstgrasser et al. (2024) showed mixed (human + synthetic) corpora substantially bound the phenomenon. The extrapolation to enterprise RAG and search-index degradation is an inference about a related but distinct mechanism (signal-to-noise dilution in mixed corpora), not a direct application of the Shumailov result. The inference is plausible and consistent with operator-reported RAG quality drift, but it is not, at the time of writing, supported by an equivalent peer-reviewed production-scale demonstration. Treat it as a working hypothesis worth instrumenting against, not as an established outcome.

11. Risk Domain 7 - Democratization of Automated Threat Capabilities

The capability uplift that LLMs provide to legitimate knowledge workers and operators applies equally - and without restriction - to threat actors. The technical barrier to executing sophisticated cyberattacks has historically been a meaningful constraint. That constraint is being systematically eroded.

11.1 LOWERED ENTRY BARRIERS FOR COMPLEX ATTACKS

Prior to widespread LLM availability, constructing a polymorphic intrusion script, generating domain-specific social engineering content, or researching target-specific vulnerability chains required substantial technical expertise and time investment. These costs functioned as natural filters: they excluded unsophisticated actors and slowed operational tempo.

LLMs reduce these friction points substantially. An actor with limited technical background can now generate functional code for web scraping, API enumeration, credential stuffing automation, or evasion techniques through iterative natural language interaction. More significantly, the production of personalized spear-phishing content - historically constrained by the time cost of target research and message crafting - is now automatable at scale. A campaign that previously required a skilled social engineer working full-time can now be partially automated, with LLMs generating target-specific narratives from publicly available information at throughput rates that human operators cannot match.

11.2 EVOLVING DEFENSIVE ECONOMICS

The asymmetry between offense and defense in this context is structural. Attackers using LLMs for content generation and reconnaissance operate with near-zero marginal cost per additional target. Defenders must evaluate each suspicious interaction individually, at full operational cost.

Static signature-based defenses - email gateways trained on prior phishing patterns, rule-based content filters, conventional IDS rulesets - are demonstrably insufficient against LLM-generated content that is syntactically novel, contextually plausible, and semantically coherent. The economics of defense have shifted: maintaining equivalent protection against AI-augmented threats requires behavioral analysis, semantic classification, and adaptive response systems that carry substantially higher operational and procurement costs than the threat they counter.

This is not a speculative future state. Security vendors including Mandiant, CrowdStrike, and Proofpoint have documented LLM-assisted threat activity in 2023–2024 operations. BunkerWeb and comparable application-layer security platforms are increasingly required to address this threat class as part of baseline WAF and behavioral filtering configuration - a requirement that was not in scope three years ago.

12. Risk Domain 8 - Traffic Intermediation and Web Hosting Consolidation

A systemic economic risk that has received insufficient technical analysis concerns the structural impact of LLM search interfaces on web traffic flows. As AI-powered answer engines - ChatGPT Search, Perplexity, Google AI Overviews, Microsoft Copilot Web Search - increasingly

serve synthesized responses to user queries, the traffic ecology of the open web is being reorganized in a way that produces asymmetric costs for content producers and infrastructure operators.

12.1 THE TRAFFIC INTERMEDIATION EFFECT

Traditional web search engines drive referral traffic: a user receives a result list, clicks a link, and arrives at the publisher's site. The publisher bears the infrastructure cost of serving that user but receives the revenue-generating visit. AI search interfaces invert this model: the system crawls, ingests, and synthesizes publisher content, then serves a generated response to the user. The user's query is resolved without a site visit.

The publisher in this model bears two costs - the bandwidth and infrastructure cost of serving the crawler that ingested the content, and the opportunity cost of the visit that no longer occurs. The economic value extracted by the AI system from the publisher's content is not redistributed to the publisher. This is a structural extraction, not a temporary side effect of a transitional technology phase.

Counter-evidence and scale check.

The strongest counter-argument is the 2024–2025 wave of publisher-AI licensing deals: OpenAI–Axel Springer, OpenAI–Associated Press, OpenAI–News Corp, OpenAI–Le Monde, OpenAI–Vox Media, OpenAI–Time, OpenAI–Reddit (~\$60M/yr), Anthropic–Reddit, and a handful of regional outlets. These deals constitute evidence that content-licensing markets are forming, and they deserve direct engagement rather than dismissal.

The pushback survives a scale check, however. Aggregate publicly-disclosed AI-licensing revenue across the publisher sector is estimated at \$150–250 M/yr as of late 2025 (sum of disclosed deal values, reported in publisher trade press). The historical publisher referral economy from organic search — the system AI-search interfaces are progressively substituting — is estimated at \$50–100 B/yr globally (Pew Research, Reuters Institute Digital News Report). Current licensing flows therefore internalise on the order of **0.2–0.5%** of the externality at issue. This is consistent with “internalisation has begun” and inconsistent with “internalisation is on a trajectory to match displaced referral value within the planning horizon of an infrastructure operator (3–5 years).” The “structural” framing is retained for that reason, with the licensing-deal evidence acknowledged as directional progress at sub-percent scale.

From an infrastructure operations standpoint, this translates into a measurable change in traffic composition: egress costs for AI crawler traffic increase, while revenue-generating human visit traffic decreases. The ratio shift is asymmetric by design. Platform operators running on pay-

per-transfer cloud infrastructure (AWS CloudFront, Cloudflare, Azure CDN) face rising bandwidth costs for content that is no longer converting into business outcomes.

12.2 MID-TIER PUBLISHER VIABILITY AND INFRASTRUCTURE CONSOLIDATION

The long-term consequence of this shift is structural consolidation. Publishers and content platforms that cannot sustain infrastructure costs without proportional traffic revenue will either exit the market, reduce content production, or migrate to paywalled or authenticated-only delivery models. Both outcomes reduce the availability of freely accessible, independently produced content on the open web.

The hosting and infrastructure layer reflects this: independent publishers running self-hosted or small-provider infrastructure face a more acute version of the economics that already pressures this segment. Mid-sized platform operators - typically the customers of regional hosting providers, colocation facilities, and managed WAF services - are the population most directly affected. Hyperscalers, by contrast, often benefit on both sides: as the providers of AI compute for the systems generating the intermediation, and as the cloud infrastructure providers capturing the remaining high-volume publisher workloads as consolidation continues.

12.3 THE RENTABILITY OF CLOUD PLATFORMS AND CROSS-SUBSIDIZATION

Beyond market consolidation, the fundamental rentability (profitability) of cloud platforms is undergoing a structural distortion. The capital expenditure (CapEx) required to construct AI-capable data centers is historically unprecedented. While hyperscalers capture new revenue streams from AI APIs, the underlying hardware-GPUs, specialized cooling, optical networking-carries massive procurement and depreciation costs, suppressing overall infrastructure margins.

To maintain the rentability of the broader cloud platform and satisfy shareholder margin expectations, operators are structurally incentivized to increase prices on standard, non-AI infrastructure. This manifests as rising costs for traditional compute instances (CPU), block storage, and egress bandwidth. The result is an invisible cross-subsidy: organizations running standard web workloads, CMS hosting, and legacy applications are effectively paying a premium to subsidize the hyperscalers' multi-billion-dollar AI infrastructure build-outs.

Regulatory gap: Current EU and US competition law frameworks were not designed to address AI-mediated traffic intermediation as a market distortion mechanism. Proposed legislation (EU AI Act, Digital Markets Act enforcement) does not directly address the extraction economics of AI crawling from independent publishers. This represents a regulatory gap with significant implications for infrastructure investment incentives across the open web.

13. Evidence-Based Recommendations

IMMEDIATE (0–30 DAYS) - TRAFFIC LAYER

Deploy behavioral traffic classification at the proxy layer. Rule-based UA matching against the documented crawler list is insufficient. Add request rate, inter-request timing, endpoint affinity (concentration on high-value content paths), and session depth as classification signals. Cloudflare, Nginx with Lua, and BunkerWeb all support custom scoring logic. Separate rate-limit buckets for declared AI crawlers, undeclared automation, and human sessions independently to avoid collateral damage.

IMMEDIATE (0–30 DAYS) - STORAGE LAYER

Audit and isolate LLM artifact directories before the next backup cycle. Identify all directories containing LLM outputs (vector stores, conversation logs, model caches, draft export folders). Apply explicit exclude rules in backup configuration for volatile, regenerable artifacts. Apply short retention policies (7–14 days) to AI intermediary outputs. Document this policy for NIS2 and GDPR record-keeping requirements.

30–60 DAYS - OBSERVABILITY

Instrument a traffic taxonomy dashboard. Without measurement, risk is unquantifiable. At minimum, report weekly: (a) share of requests by traffic class (human / known AI crawler / unclassified automated / security scanner), (b) storage growth rate segmented by AI-artifact directories vs. business data, (c) backup job duration trend, (d) SIEM event ingestion rate vs. capacity limit. These four metrics provide early warning across both risk domains.

30–60 DAYS - COMPLIANCE

Extend GDPR data inventory to AI artifact types. If your organization uses any LLM tool that processes user-provided content or web sessions, that tool's output logs may contain personal data. Per GDPR Article 30, these must appear in your Record of Processing Activities. Apply storage limitation under Article 5(1)(e) explicitly. Under the EU AI Act, if any AI system in use qualifies as high-risk under Annex III, ensure log retention meets Article 12 technical standards - structured, traceable, time-bounded.

30–60 DAYS - GOVERNANCE

Audit promotional posture as a deliberate AI-exposure control. Per the per-site heterogeneity in §4.4, AI-training crawler pressure is jointly determined by content profile *and* public discoverability. Promotional posture (sitemap submission, structured-data markup, advertising-driven inbound links, presence in directories the AI seed graph traverses) is therefore a tunable surface that is distinct from technical opt-out (`robots.txt` , AI-bot UA blocks) and composes multiplicatively with it. For internal applications, staff portals, and properties whose business value does not depend on third-party search referral, deliberately limiting promotional posture can reduce AI-training pressure by 1–2 orders of magnitude (as observed between Sites A and C in the fleet) without any technical crawler-blocking. For revenue-bearing public properties, the lever cannot be used wholesale, but it should be evaluated per-property rather than applied as a single site-wide setting. Cost is process / governance rather than capex.

60–90 DAYS - ARCHITECTURE

Implement lifecycle management on all cloud sync services. Microsoft 365 administrators can configure retention labels, auto-delete policies, and sensitivity labels through Microsoft Purview. Google Workspace administrators can configure retention rules in Google Vault. Both support policy-based deletion of content meeting defined criteria. Apply these to AI output folders explicitly, with documented justification. Test OneDrive and Google Drive quotas against projected AI output volume growth quarterly.

90 DAYS+ - CAPACITY PLANNING

Decouple storage capacity planning from headcount-linear assumptions. Traditional storage forecasting assumes storage grows with headcount and business volume. LLM workloads break this assumption: a single AI deployment can generate data volumes equivalent to dozens of additional human users. Establish a separate AI workload storage budget, with quarterly review cadence tied to AI tool adoption metrics - not just headcount.

14. Open Questions and Research Gaps

Several dimensions of this risk landscape remain under-researched or undisclosed:

- **Cumulative crawl volume per site:** No major AI company has disclosed total crawl volume per target domain. The operational impact on individual sites - particularly small and mid-sized publishers - remains largely unmeasured in peer-reviewed literature.

- **Energy and carbon footprint of redundant AI artifact storage:** The environmental cost of storing, replicating, and backing up LLM-generated artifacts that are never retrieved is not yet quantified at industry scale.
- **Insurance and liability:** Standard cyber insurance policies were not written with AI crawler DDoS, AI-artifact storage failures, or prompt-injection via web content in mind. Coverage gaps are not yet adjudicated.
- **Aggregate training data duplication:** Multiple AI companies independently crawling the same content creates global-scale data redundancy. The infrastructure cost of this duplication - bandwidth, storage, processing - is not publicly accounted.

7. Risk Domain 3 - Cognitive Asymmetry and Human Exhaustion

An often-overlooked consequence of machine-scale AI operations is the cognitive toll on the human operators and end-users on the receiving end. The asymmetry between the zero-cost generation of AI traffic and the high-cost human triage required to manage its fallout creates structural exhaustion across three distinct personas:

- **The Independent Operator / Hobbyist:** Individuals maintaining personal servers, self-hosted services, or small community nodes are ill-equipped for this wave. Lacking enterprise WAF orchestration, they often face sudden resource exhaustion. An individual operator may wake up to a crashed server because a new undeclared AI crawler decided to index their entire historically generated photo gallery in a single hour.
- **The SOC / CSIRT Analyst:** Security teams are drowning in a "Signal-to-Noise Ratio" crisis. When AI agents generate thousands of anomalous HTTP requests mimicking rapid traversal or vulnerability scanning, standard SIEM alerts trigger continuously. Evaluating whether an IP behaving erratically is a malicious actor or just a misconfigured LangChain web-scraping script consumes analyst bandwidth, leading to alert fatigue and the genuine risk of missing directed, human-driven attacks.
- **The End User / Knowledge Worker:** Downstream users are equally affected by the output bloat. The volume of "AI-enhanced" summaries, auto-generated reports, and endless versioning litters collaboration tools like SharePoint or Teams. Information retrieval becomes harder when search results are diluted by AI-generated noise, creating daily micro-frictions as users struggle to locate authentic artifacts.

7.1 REAL-WORLD PATHOLOGY: PRESTASHOP TELEMETRY EXHAUSTION

The cognitive and infrastructural exhaustion is acutely visible on legacy architectures not built for infinite artificial traversal. A documented example is the e-commerce platform PrestaShop. By design, native PrestaShop instances track visitor statistics directly inside the relational database (via the `ps_connections`, `ps_guest`, and `ps_page_viewed` tables) rather than relying exclusively on flat access logs.

This is not a marginal platform effect in France: the 2026 Friends of Presta barometer (published by E-Commerce Nation) reports PrestaShop at 19.3% of active e-commerce sites (24,211 sites), while also leading by cumulative revenue at EUR 7.96 billion. In operational terms, this means telemetry-related failure modes on PrestaShop affect a material share of real-world commerce rather than a niche technical segment.

This exposure also includes a long tail of amateur and semi-professional operators who rely on PrestaShop for niche catalog commerce, including hobbyist ecosystems such as 3D-printed figurines, tabletop accessories, maker components, and small-batch collectible merchandise. These operators typically lack dedicated SRE capacity, making them disproportionately vulnerable to alert overload, database bloat, and observability blind spots when crawler pressure rises.

For amateur, semi-professional, and professional merchants alike, business continuity depends on the shop staying fully responsive. If the storefront slows down or fails, users abandon sessions, conversion drops immediately, and revenue is lost in real time. The cognitive burden then shifts to shop owners and their informal IT support network (friends, freelancers, or part-time admins), who are often forced to troubleshoot outages without clear root-cause visibility and without a deep understanding of why the platform is degrading under automated traffic pressure.

When subjected to multi-threaded LLM crawling, this architecture becomes catastrophic. A swarm of AI agents extracting product data generates an immediate explosion of rows in these tracking tables. An administrator expecting to analyze human customer journeys is instead confronted with gigabytes of database bloat. The database grows to the point where standard cron-based optimization scripts time out. Administrator dashboards freeze trying to render statistics, effectively blinding the site owner to real commercial activity while silently pushing the underlying MySQL/MariaDB server to its I/O limits.

9. Risk Domain 5 - Cognitive Accessibility and Vulnerable User Interaction

While the cognitive risks for human operators and young users are documented in the preceding sections, a distinct and clinically significant risk dimension applies to adult users with pre-existing psychological vulnerabilities, neurodivergent profiles, or social accessibility deficits. The architecture of conversational AI systems - engineered for engagement, continuity, and frictionless interaction - creates structural conditions that may systematically disadvantage these populations.

Methodological note: Longitudinal peer-reviewed research on LLM-specific interaction effects on vulnerable adult populations remains limited. The patterns documented below are extrapolated from established research in technology addiction, human-computer interaction, and parasocial relationship formation. They represent risk hypotheses grounded in established behavioral models, not confirmed outcomes.

9.1 ASYMMETRIC SOCIAL DYNAMICS

Individuals with social anxiety disorders, autism spectrum conditions, or social communication differences often find that the low-friction, non-judgmental interaction architecture of conversational AI systems provides immediate relief from interpersonal costs. Unlike human interlocutors, LLMs do not demonstrate impatience, shift topics unexpectedly, or impose conversational norms that require real-time social processing.

From an accessibility perspective, this is a documented benefit. From a risk perspective, it is also a pathway to substitution: when an AI system reliably provides perceived social connection with zero interpersonal cost, it may progressively displace the effortful, unpredictable, but developmentally essential experience of human social interaction. This substitution risk is structurally invisible to the system, which has no mechanism to distinguish therapeutic interaction from pathological dependence - and no incentive to do so.

9.2 EPISTEMIC OVER-RELIANCE AND MOTIVATED VALIDATION

LLMs respond to prompts as stated. They do not diagnose the premise. A user experiencing health anxiety who asks "what are the symptoms of [condition]?" will receive a detailed, authoritative-sounding answer. The system will not probe whether the question reflects genuine clinical concern, hypochondriacal preoccupation, or a misframing of the actual problem.

This creates a structurally asymmetric epistemic environment: users who present incorrect or anxious framings receive confident, detailed responses that validate the framing by engaging with it. Over repeated interactions, this can reinforce pre-existing cognitive distortions - a pattern well-documented in research on confirmation bias and availability heuristic amplification through digital media, now extended to an interactive, personalized, high-verbosity medium.

9.3 UNSTRUCTURED HEALTH AND QUASI-THERAPEUTIC INTERACTIONS

A significant and growing subset of LLM use occurs in quasi-therapeutic contexts: users discussing personal distress, suicidal ideation, relationship crises, or mental health symptoms with AI systems. Unlike regulated mental health platforms, general-purpose LLMs operate without clinical oversight, crisis detection protocols, or escalation pathways.

This gap has infrastructure implications. When a platform inadvertently becomes a crisis intervention point - without the engineering, training, or regulatory compliance of clinical systems - it assumes a risk liability that is neither scoped nor disclosed. The failure mode is not theoretical: documented cases exist of AI systems providing factually incorrect, emotionally reinforcing, or inappropriately permissive responses to users in acute distress. From a compliance standpoint, the EU AI Act's classification of high-risk AI systems under Annex III specifically includes systems used in safety-critical decision contexts - a framing that may extend to health-adjacent conversational AI as regulatory interpretation matures.

9.4 PERSISTENT ENGAGEMENT LOOPS AND EXECUTIVE FUNCTIONING

LLM interfaces are architecturally unbounded. There are no natural session-termination signals equivalent to the end of a book chapter, the conclusion of a video, or the fatigue of a human interlocutor. This infinite generation architecture may pose particular risk for users with conditions affecting executive functioning, impulse regulation, or time estimation - including ADHD, bipolar spectrum conditions, and certain anxiety disorders.

The combination of on-demand responsiveness, high information density, and absence of natural stopping points creates persistent engagement loops with no equivalent in prior media. This is not a feature that requires exploitation or adversarial engineering - it is the default operating condition of the system.

7.2 A GENERATION GROWING UP INSIDE AN UNCONTROLLED EXPERIMENT

Epistemic status: The mechanisms below draw on established pre-LLM screen research and developmental psychology. The specific effects of LLM-era AI interaction on youth cognition are *not yet longitudinally studied*. This section therefore distinguishes documented evidence from qualified research gaps. The absence of data is itself a risk indicator.

The cognitive risks described in this article do not spare minors - and in their case, the unknowns are far more profound. Societies are deploying LLM systems at population scale without longitudinal evidence of how persistent, interactive AI exposure affects developing cognition. We are, in effect, conducting an uncontrolled experiment on children with no control group and no mechanism for informed consent.

7.2.1 WHAT SCREEN AND INTERNET RESEARCH ACTUALLY TELLS US

Existing research on screens and internet exposure was largely conducted before the LLM era. Key findings include:

- **Screen time and adolescent mental health (Twenge, 2017–2023):** Longitudinal data across multiple cohorts shows a statistically significant correlation between increased screen time - specifically smartphone and social media use - and elevated rates of anxiety, depression, and loneliness among adolescents aged 12–17, particularly girls. This correlation accelerated after 2012 (peak smartphone adoption). Jean Twenge's research, spanning 11 million participants over multiple decades, documents measurable divergence in adolescent mental health trajectories coinciding with internet adoption patterns.
- **Haidt's "Anxious Generation" (2024):** In his 2024 book, Jonathan Haidt synthesizes epidemiological, psychological, and sociological data to argue that the combination of smartphone adoption and social media use during early adolescence is causally associated with the widespread mental health deterioration observed across North America, Europe, and Australia since 2012. While causality debates continue in the academic community, the temporal correlation and cross-national consistency of the data are considered significant. *Critically, all this research predates the LLM era by at least a decade.*
- **PISA 2022 - Declining Reading Comprehension:** OECD's Programme for International Student Assessment (PISA 2022) recorded the largest cross-national decline in reading comprehension scores since the programme's inception. 15-year-olds in the majority of measured countries showed deterioration that pre-pandemic baselines cannot fully explain. Researchers note temporal alignment with digital media saturation, though causality has not been established with certainty.

- **WHO screen time guidelines (2019):** The World Health Organization recommends no screen time for children under 2, maximum one hour per day for ages 2–5. These guidelines were developed without data on AI-mediated interaction and do not account for conversational AI systems that behaviorally differ from passive video consumption.

7.2.2 WHY LLM-ERA EXPOSURE IS QUALITATIVELY DIFFERENT

All prior research concerns passive or broadcast-style digital media: video, social feeds, search engines. LLMs introduce a categorically new dynamic - the system responds. It adapts. It provides on-demand answers that feel authoritative. This creates several vectors of concern that existing research does not address:

- **Metacognitive offloading:** When a child outsources reasoning to an AI that produces confident, readable, plausible-sounding outputs, the cognitive work of forming a judgment - evaluating sources, tolerating ambiguity, sitting with uncertainty - is no longer performed. Whether persistent offloading inhibits the development of autonomous critical thinking capacity is not yet empirically studied. It is a legitimate and measurable research question. We do not yet have the answer.
- **Epistemic confusion:** Adults struggle to distinguish AI-generated from human-authored text. Children and adolescents, with less world knowledge and fewer heuristics, are more exposed. Growing up in an information environment where authoritative-sounding text may or may not be grounded in reality - and where the generation mechanism is invisible - represents a developmental condition with no historical precedent.
- **Parasocial dependency:** Conversational AI systems are engineered for engagement. They do not tire, judge, or reject. For younger users - particularly those already socially isolated or anxious - the risk of forming affectively asymmetric dependencies (where the user ascribes emotional meaning to an AI interaction) is real. Unlike classical parasocial relationships with celebrities, AI systems respond and adapt, creating a qualitatively more immersive dynamic.
- **Attention architecture:** LLM-powered interfaces are designed to produce complete answers, reducing the need for exploratory search, reading, synthesis, and conclusion-forming. The long-form reading and inference skills tracked by PISA - already declining - may face further pressure from a generation that grows up with access to a system that does the synthesis for them.

7.2.3 THE INFRASTRUCTURE-LEVEL RISK

From a systemic infrastructure risk perspective, this translates into a long-horizon human capital concern: the pipeline of future engineers, analysts, and operators capable of understanding, maintaining, and securing complex digital infrastructure depends on a generation developing the relevant cognitive skills. If LLM adoption at the educational level accelerates metacognitive offloading during formative years, the talent pipeline for infrastructure operations is exposed to a structural risk that will not manifest until the 2030s - but that begins accumulating now.

There is also a more immediate political risk. Populations that cannot distinguish AI-generated information from primary reporting, and that have been exposed since childhood to systems that confidently answer any question, are more susceptible to coordinated influence operations at scale. Infrastructure defense requires human operators who think adversarially, skeptically, and laterally - traits associated with high tolerance for ambiguity and comfort with incomplete information. These traits are shaped in part during adolescence. We do not yet know whether growing up with AI tutors shapes or erodes them.

What can be said with precision is this: **we do not know**. We do not have the data. The absence of longitudinal research on LLM-era cognitive development is not reassuring - it is itself a risk signal. Societies and infrastructure organizations have a reasonable basis to apply the precautionary principle: acknowledge the knowledge gap explicitly, fund independent longitudinal research, and avoid treating the absence of confirmed harm as evidence of safety.

15. Conclusion

The hidden costs of LLM-scale automation are already present in production telemetry, and they are unevenly distributed. The eight risk domains catalogued in this Black Paper do not all share the same incidence pattern — and the unifying “externality” framing requires the following two-track distinction to remain defensible:

CROSS-ORGANISATIONAL EXTERNALITIES

Cost-bearer and load-generator are distinct entities. Mitigation requires either market mechanisms (content licensing), policy (mandatory disclosure, fair-compensation rules), or perimeter defence (WAF, rate-limit, robots enforcement, promotional-posture management).

Applies to: third-party publisher crawl load (§4), public energy/water/semiconductor displacement (§8), harm to vulnerable users (§9), mid-tier publisher pressure (§12).

INTRA-ORGANISATIONAL TRADE-OFFS

The AI-adopting organisation is both the load-generator and the cost-bearer. Mitigation is a governance and operational-discipline matter: lifecycle policy, baseline instrumentation, capacity planning.

Applies to: AI artifact storage growth on the adopter's own cloud (§5), enterprise RAG / knowledge-base contamination (§10.2), SIEM volume growth on the adopter's own pipeline (§6), operator cognitive load (§7).

Both tracks are real, both are measurable today, and both are visible in the fleet telemetry presented in §4.4 and Annex A. The operational implication is that AI infrastructure governance is not a single problem with a single response: cross-organisational risks demand engagement with markets and regulators in addition to perimeter defence, while intra-organisational risks demand internal lifecycle discipline that an external regulator cannot impose. Conflating the two produces either misallocated regulatory attention or misallocated engineering budget.

A third dimension surfaces from the multi-site fleet view (§4.4) that is not typically named in the AI-infrastructure literature: **discoverability** — whether a property is reachable by the AI crawler seed graph at all — is a control surface distinct from both technical opt-out (`robots.txt`) and content profile. The fleet shows AI-training pressure varying by more than two orders of magnitude between sites of comparable WAF posture, with promotional intensity (SEO, advertising, sitemap submission, inbound link campaigns) as the most plausible explanatory variable beyond content type. For operators whose property value does not depend on third-party search referral, promotional posture is a tunable lever that has been overlooked. For operators whose property value does depend on it, the lever cannot be used wholesale — but it can be applied per-property, which is a finer-grained governance question than the field currently asks.

Where this Black Paper stops short on purpose: it does not attempt a comparison framework against the other 2026 infrastructure-attention competitors (ransomware-as-a-service evolution, post-quantum-crypto migration, cloud-concentration risk, supply-chain compromise, DORA/CRA regulatory shifts). Without that comparison, this document should not be read as a claim that AI infrastructure risk is the top-priority 2026 concern — only that it is a sufficiently

material concern, with sufficiently identifiable incidence patterns, to merit dedicated instrumentation and governance work. The companion White Paper (in draft) will provide the comparison framework alongside the mitigation playbooks.

Disclosure: the field observations in §4.4 and Annex A were collected from BunkerWeb-protected production sites operated by the author. The recommendations name BunkerWeb among other reverse-proxy and WAF options (Cloudflare, Nginx-with-Lua); the author has no commercial relationship with the BunkerWeb project beyond operating it as a user. The fleet harvest tooling used to produce §4.4's aggregates is open-source and reproducible (`harvest.report` , MIT, schema `bw.harvest.v3`).

Annex A. Verified Field Telemetry (Anonymized Site)

Anonymization note: The production domain, brand, and category labels were removed. The dataset below is presented as *Site A* to prevent direct targeting while preserving operational signal.

This annex embeds telemetry extracted from consolidated reverse-proxy and WAF access logs for an anonymized e-commerce workload (**Site A**) over a 17-day observation window (26-Apr-2026 to 12-May-2026). Data integrity checks were performed before integration: daily aggregates were recomputed and verified against the global totals, with exact equality on request counts, bytes transferred, and blocked-request counters.

VERIFIED AGGREGATE RESULTS (SITE A)

METRIC	VALUE	INTERPRETATION
Total requests	8,697,962	High-volume perimeter pressure in less than three weeks
AI-classified requests	7,153,371 (82.24%)	Automation dominates traffic composition
Traditional bots	745,962 (8.58%)	Classical crawlers remain significant but secondary
Human traffic	798,629 (9.18%)	Human share is structurally compressed
Total transferred bytes	920,369,355,879	~920.37 GB served during the observed period
AI byte share	878,038,133,231 (95.40%)	Bandwidth burden is overwhelmingly AI-driven
Blocked AI requests (HTTP 403)	1,036,427 (14.49% of AI requests)	Protection controls engage at sustained high rates
Category-page traversals	2,482,198 total; 1,947,214 AI (78.45%)	Deep catalog traversal is mostly machine-driven

A.1 OPERATIONAL READING

- **Asymmetric load:** AI traffic is not only dominant in request count but disproportionately dominant in bytes served, confirming that the highest infrastructure cost center is machine-origin demand.
- **Defense pressure:** A seven-figure count of blocked AI requests in 17 days indicates sustained adversarial or at least non-cooperative automation pressure at the edge.
- **Catalog focus:** Most high-frequency automation targets category/listing navigation paths, not only top-level landing pages, which amplifies backend query and cache-miss costs on dynamic commerce stacks.
- **Continuity exposure:** At this magnitude, proxy/WAF, logging, SIEM ingestion, and backup systems become coupled risk surfaces rather than independent layers.

A.2 METHODOLOGICAL NOTE

Classification used deterministic User-Agent families (AI crawlers, traditional bots, residual human traffic) plus status-code distribution and URL-pattern counters. The annex intentionally excludes raw domains, full URL labels, and direct commercial identifiers. The objective is reproducible risk characterization without publishing targetable infrastructure fingerprints.

Annex B. Reader-Reproducible Discoverability Audit

The §4.4 fleet observation and the §13 promotional-posture recommendation both rest on the claim that AI-training crawler pressure correlates with public discoverability, not only with content profile or technical opt-out. The check below allows any operator with shell access to a property they control to produce a first-order discoverability signal for that property, in under five minutes, without privileged third-party data. It is not a substitute for paid SEO or referer-graph audits; it is a lower-bound observational baseline.

Scope. The audit covers four signals: (1) sitemap presence and URL count; (2) `robots.txt` directives for AI crawlers; (3) presence in Common Crawl indexed-URL counts (sampling, not exhaustive); (4) a normalised promotional-posture score combining the prior three. It does not measure inbound link graph, ad-spend, or third-party directory presence; those require paid data sources.

B.1 AUDIT SCRIPT

Save the following as `discoverability-audit.sh`, make executable (`chmod +x`), and invoke as `./discoverability-audit.sh https://your-property.example`. Requires `curl`, `grep`, and `wc` (BusyBox-compatible).

```

#!/usr/bin/env bash
# discoverability-audit.sh - first-order AI-discoverability signal
# Usage: ./discoverability-audit.sh https://your-property.example
set -euo pipefail
URL="${1:-}"
if [[ -z "$URL" ]]; then echo "Usage: $0 https://your-property.example" >&2; exit 2;
fi
HOST="$(echo "$URL" | sed -E 's#^https?://([^/]+).*#\1#')"
echo "=== Discoverability audit: $HOST ==="

# 1. Sitemap presence + URL count
echo "--- 1. Sitemap ---"
for SM in sitemap.xml sitemap_index.xml sitemap-index.xml; do
    CODE="$(curl -s -o /tmp/sm.$$ -w '%{http_code}' "$URL/$SM" || echo 000)"
    if [[ "$CODE" == "200" ]]; then
        COUNT="$(grep -c '<loc>' /tmp/sm.$$ || echo 0)"
        echo "  $SM: HTTP 200, ${COUNT} <loc> entries"
    fi
done
rm -f /tmp/sm.$$

# 2. robots.txt AI directives
echo "--- 2. robots.txt AI directives ---"
curl -s "$URL/robots.txt" -o /tmp/rb.$$ || echo " (no robots.txt)"
if [[ -s /tmp/rb.$$ ]]; then
    for UA in GPTBot ChatGPT-User ClaudeBot Claude-Web anthropic-ai Google-Extended
CCBot PerplexityBot meta-externalagent FacebookBot Bytespider; do
        if grep -qi "User-agent:.*$UA" /tmp/rb.$$; then
            echo "  $UA: declared"
        fi
    done
fi
rm -f /tmp/rb.$$

# 3. Common Crawl presence (sample - latest monthly index)
echo "--- 3. Common Crawl presence (sample) ---"

```

```

CC_INDEX="$(curl -s https://index.commoncrawl.org/collinfo.json | grep -oE '\"cdx-
api\":"[^\"]+\'' | head -1 | sed 's/\\"cdx-api\":"//;s/\\/\\/')"
if [[ -n "$CC_INDEX" ]]; then
    CC_COUNT="$(curl -s "${CC_INDEX}?url=${HOST}/*&output=json&limit=1000" | wc -l)"
    echo " Latest monthly index: ${CC_COUNT} URLs indexed (capped at 1000 sample)"
else
    echo " (Common Crawl index unreachable)"
fi

echo "--- Done ---"
echo "Interpretation:"
echo " - High sitemap count + few robots blocks + high CC presence => HIGH
discoverability"
echo " - No sitemap or AI-bot blocks declared + low CC presence => LOW
discoverability"
echo " - Compare across your fleet; flag outliers per direction."

```

B.2 INTERPRETATION GUIDE

- **Sitemap density:** A property with >10,000 sitemap URLs and no AI-bot `robots.txt` exclusions is operating at the high end of the discoverability spectrum. A property with no sitemap and 5+ AI-bot exclusions is at the low end.
- **Common Crawl presence:** Properties appearing in Common Crawl's latest monthly index with non-trivial URL counts (>100 sample-capped) are in the seed graph that most frontier AI training pipelines downstream from. Absence from Common Crawl is not proof of low pressure (other crawler graphs exist) but presence is strong evidence of high pressure.
- **Cross-fleet comparison:** The audit's primary value is differential. Run it across an operator's full property set; the property that ranks highest by sitemap-count×CC-presence and lowest by robots-exclusion-count is the property most exposed to AI-training crawler load, all other factors equal.
- **Limits:** No measure of inbound link graph, advertising spend, or paid-directory inclusion. Those require third-party data (Ahrefs, Semrush, SimilarWeb) and are out of scope for a free, reader-reproducible check.

References

1. Imperva. *Bad Bot Report 2024*. Imperva Research Labs, April 2024. Available at imperva.com/resources/resource-library/reports/bad-bot-report/
2. OpenAI. *GPTBot documentation*. August 2023. Available at platform.openai.com/docs/gptbot
3. Google Search Central. *Google-Extended control for AI model training*. September 2023. Available at developers.google.com/search/docs/crawling-indexing/google-extended
4. Common Crawl Foundation. *Common Crawl Statistics and Data Overview*. commoncrawl.org
5. IDC. *The Digitization of the World - From Edge to Core (Data Age 2025)*. Seagate-sponsored IDC white paper, November 2018, with subsequent updates 2022–2024.
6. Reinsel, D., Gantz, J., Rydning, J. *The Digitization of the World*. IDC White Paper, 2018. doc-number US44413318.
7. European Parliament and Council. *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (EU AI Act)*. Official Journal of the European Union, July 2024.
8. European Parliament and Council. *Regulation (EU) 2016/679 (GDPR)*. Official Journal of the European Union, May 2016.
9. European Parliament and Council. *Directive (EU) 2022/2555 on measures for a high common level of cybersecurity (NIS2)*. Official Journal of the European Union, December 2022.
10. Cloudflare. *Cloudflare Radar - Bot Traffic Trends*. radar.cloudflare.com (continuously updated, 2024 data referenced).
11. Akamai Technologies. *State of the Internet: Security Report 2024*. akamai.com/resources/state-of-the-internet-report
12. Perez, E., et al. *Ignore Previous Prompt: Attack Techniques For Language Models*. NeurIPS 2022 Workshop on Machine Learning Safety. (Prompt injection foundational research.)
13. Shumailov, I., Shumaylov, Z., Zhao, Y., Gal, Y., Papernot, N., Anderson, R. *The Curse of Recursion: Training on Generated Data Makes Models Forget*. Nature, July 2024.
14. Gerstgrasser, M., Schaeffer, R., Dey, A., Rafailov, R., et al. *Is Model Collapse Inevitable? Breaking the Curse of Recursion by Accumulating Real and Synthetic Data*. arXiv:2404.01413, April 2024.
15. Anthropic. *Claude model card and usage policies*. anthropic.com/model-card (referenced for ClaudeBot documentation).
16. Microsoft. *FY2024 Annual Report*. microsoft.com/en-us/investor/annual-reports.aspx
17. Microsoft. *OneDrive for Business storage policies*. docs.microsoft.com (versioning and retention capabilities).
18. ISO/IEC 27001:2022. *Information security management systems - Requirements*. International Organization for Standardization.
19. E-Commerce Nation / Friends of Presta. *Barometre CMS e-commerce en France : Shopify domine les creations, PrestaShop le chiffre d'affaires*. March 2026. Available at ecommerce-nation.fr/barometre-cms-ecommerce-shopify-creations-prestashop-chiffre-affaires/
20. International Energy Agency (IEA). *Electricity 2024: Analysis and forecast to 2027*. IEA Publications, 2024. (Referenced for global electricity carbon intensity data.)

21. Strubell, E., Ganesh, A., McCallum, A. *Energy and Policy Considerations for Deep Learning in NLP*. Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL), 2019. (Foundational research on LLM training energy consumption.)
22. OpenAI. *Inference capacity disclosure and energy usage trends*. Internal disclosures via annual reports and blog posts, 2023–2025. (Referenced for current inference workload estimates.)
23. Nvidia. *GPU Demand and Supply Chain Analysis*. Investor Relations and market reports, 2023–2025. (Referenced for GPU H100/H200 scarcity.)
24. Patterson, D., et al. *The Carbon Footprint of Machine Learning Training Will Plateau, Then Shrink*. Computer, IEEE, 2021. (Research on LLM carbon accounting and mitigation.)
25. Luccioni, A.S., Mahendran, A. *Quantifying the Carbon Emissions of Machine Learning*. arXiv:1910.09700, 2019. (Methodology for carbon footprint estimation in AI.)
26. Twenge, J.M., et al. *Increases in Depressive Symptoms, Suicide-Related Outcomes, and Suicide Rates Among U.S. Adolescents After 2010 and Links to Increased New Media Screen Time*. Clinical Psychological Science, 2018. (Longitudinal correlation between screen time and adolescent mental health.)
27. Haidt, J. *The Anxious Generation: How the Great Rewiring of Childhood Is Causing an Epidemic of Mental Illness*. Penguin Press, March 2024.
28. OECD. *PISA 2022 Results (Volume I): The State of Learning and Equity in Education*. OECD Publishing, 2023. (Documents largest recorded cross-national decline in reading comprehension scores.)
29. World Health Organization. *Guidelines on Physical Activity, Sedentary Behaviour and Sleep for Children under 5 Years of Age*. WHO Press, 2019. (Screen time guidelines, predating LLM-era AI interaction.)

Black Paper - Infrastructure Risk Series

This document presents the adversarial, risk-focused perspective on LLM deployment at global scale. It is the first of a two-part series. The companion White Paper (constructive pathways, mitigations, opportunities) is currently in draft. All factual claims sourced from publicly verifiable primary sources. Directional projections qualified as estimates. No proprietary data used. No AI-generated statistics presented as factual without qualification.

Published: May 2026 · Author: Bryce SIMON · Co-author: Ifrit (AI) · License: CC BY 4.0 · Companion White Paper: in draft