The DPG Defence Playbook

Navigating the Paradox of Open to Mitigate AI Exploitation

Digital Public Goods Alliance

DPG

The DPG Defence Playbook

Navigating the Paradox of Open to Mitigate AI Exploitation

Authors

Lea Gimpel, Bolaji Ayodeji

Published by **Digital Public Goods Alliance Secretariat**

Version 1.0 I July 2026

Disclaimer

This document was drafted and edited by the authors with assistance from generative AI for language refinement and structural feedback. All arguments, sources, and factual claims were reviewed and verified by the author and editorial team.

Contents

Introduction

The Challenges Faced by DPG Product Owners

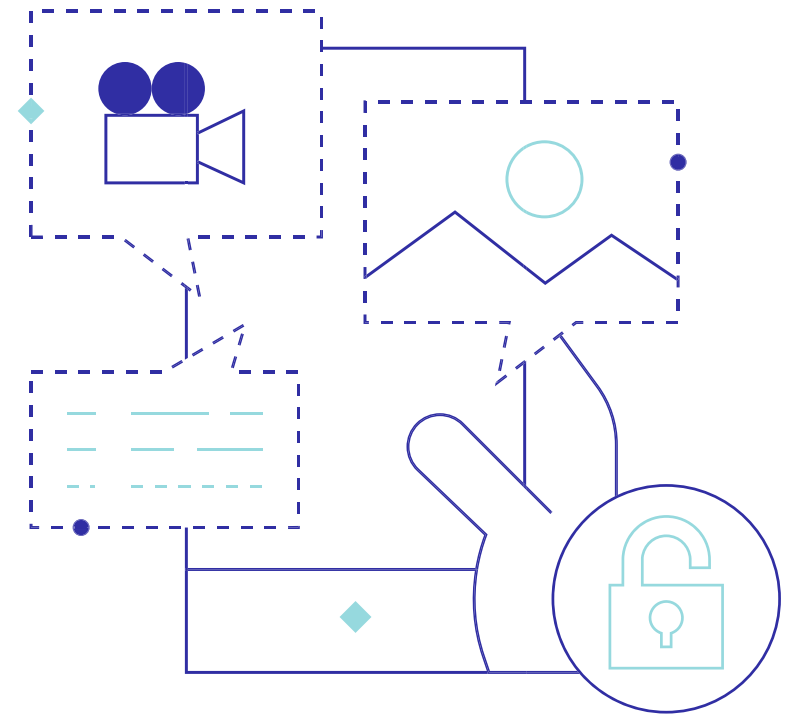

The open web is facing a systemic structural transformation. Digital public goods (DPGs) and the broader knowledge commons are increasingly vulnerable to blunt-force, automated extraction by AI crawlers and scrapers operating without reciprocity, breaking the implicit social contract of the open web and jeopardising their survival.

For open data and content DPGs, many of which rely on community- and voluntarily-driven maintenance efforts, this extraction introduces severe challenges:

Ballooning Infrastructure Costs

Automated scraping strains resources, causing bandwidth spikes and complicating content caching, impacting the ease with which humans can access open content and a project's sustainability.

Loss of Human Traffic

As generative AI models increasingly serve users with "zero-click" summaries directly on third-party platforms, human referral traffic to the original sites is dropping drastically. This lack of visible human engagement impacts downstream metrics, making traditional traffic-reliant funding and monetisation models difficult to sustain.

Erosion of Human-Curated Information

Self-sustaining editorial pipelines and communities that create trusted ground truth, such as Wikipedia, see a decrease in human engagement, while machines flood the open web with low-quality content ("AI slop"), which undermines trust in information ecosystems.

Perpetuating Power Imbalances

Those with the most resources scrap the most without giving back, denying communities the benefits of their work. This dynamic especially impacts those historically disadvantaged.

The Paradox of Openness

DPG product owners and stewards of open knowledge projects are increasingly forced to adopt measures to defend their infrastructure and data to counteract these challenges. This creates an "**openness paradox**": a tension between open access to knowledge and protecting an open resource's survival through legal and technical barriers, thereby limiting others' ability to use it without restrictions.

Keeping DPGs, including digital commons, open is crucial for several reasons:

Trust and Democracy

Human-curated open data and content counters misinformation and AI "slop" while providing an ethical, publicly accountable data foundation, including for future AI development.

Democratised Innovation

Keeping information open removes restrictive paywalls, boosting public-sector efficiency and fueling private-sector innovation.

Preserved Heritage

The ecosystem protects cultural and linguistic diversity, ensuring localised knowledge and collective heritage remain accessible rather than locked away.

Inclusion and Resilience

Open data makes marginalised communities and underrepresented geographies visible, while decentralised community governance ensures these systems remain sustainable.

Global Collaboration

DPGs and other open projects supercharge global collaboration, enabling researchers and governments to build shared knowledge bases and tools, e.g. for crisis response.

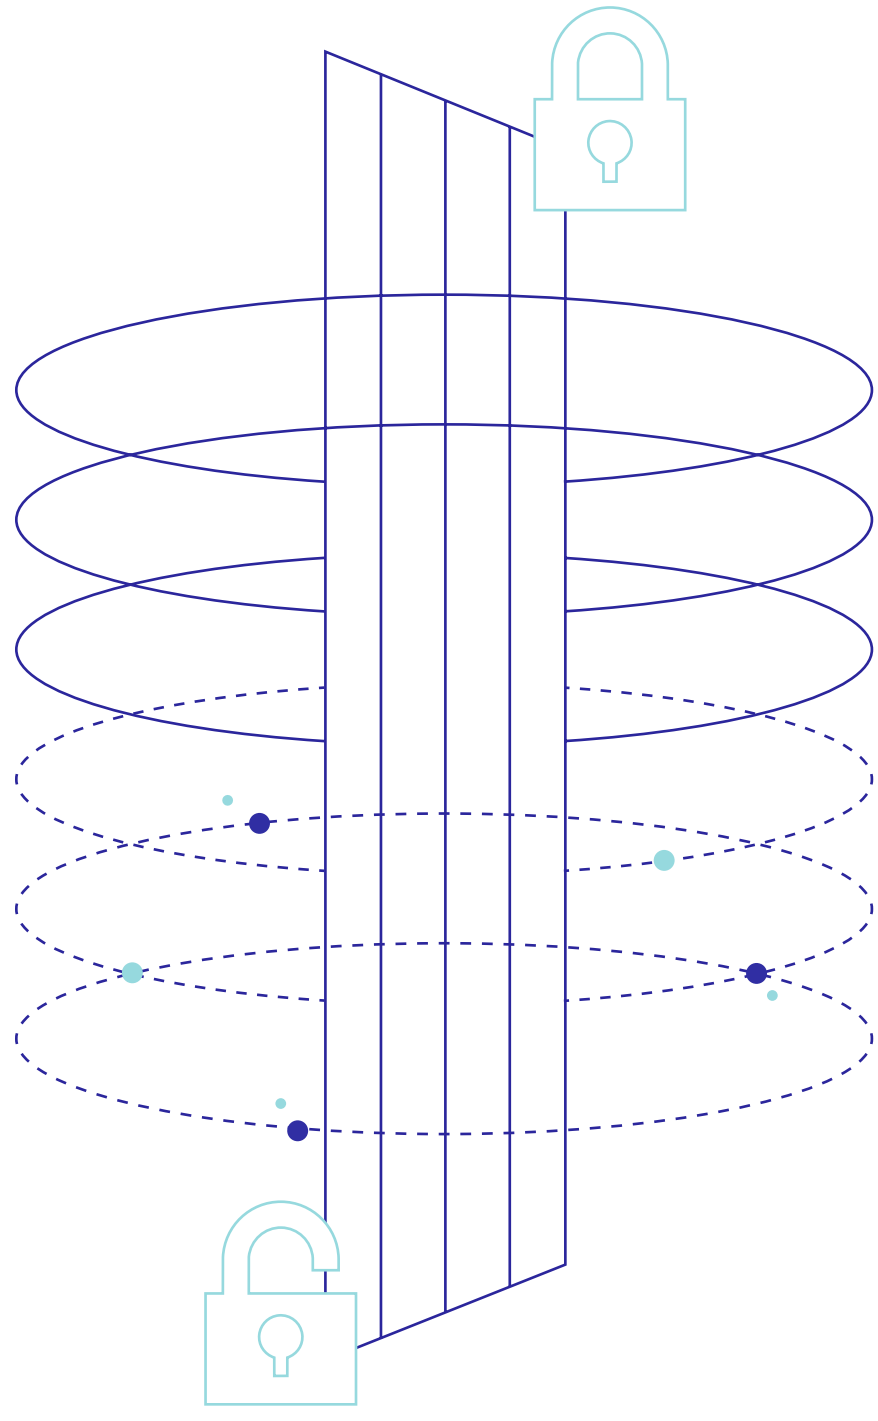

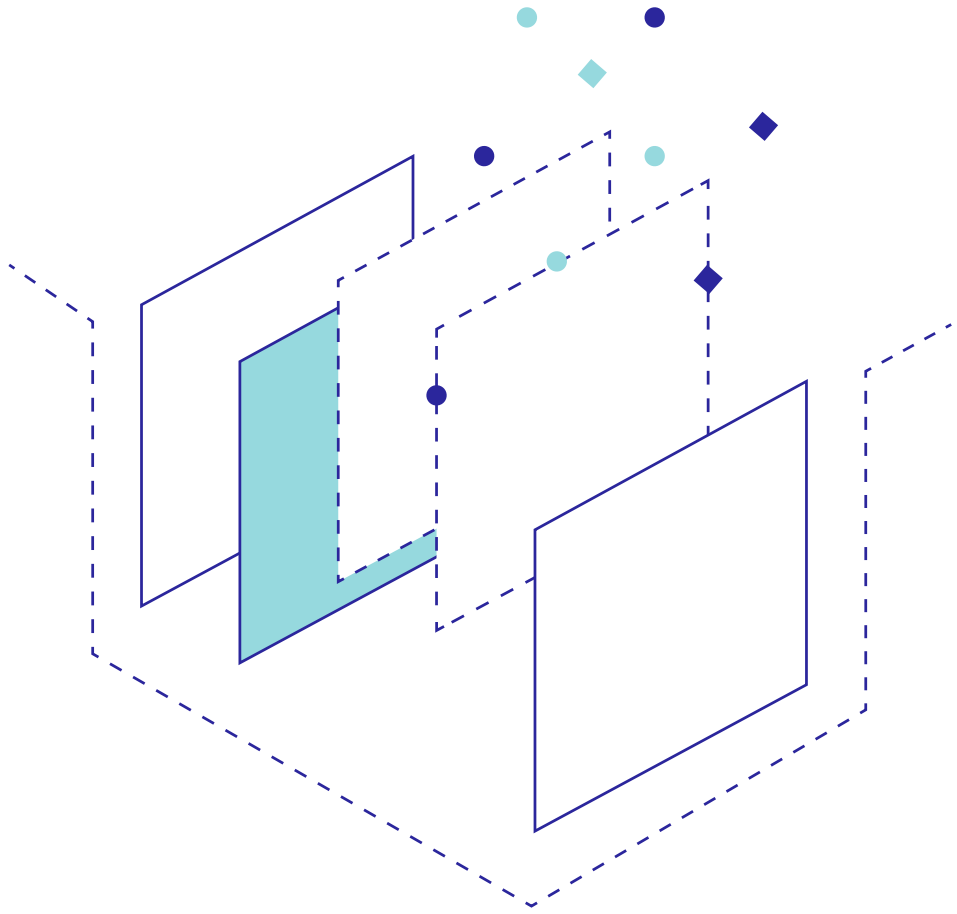

To navigate the tension between open access and protecting the resource, DPG product owners have defined best practices for AI's use of open data and content. In practical terms, DPGs must operate under the principle of being "as restrictive as necessary, as open as possible," and ensure that they meet the DPG Standard.

The DPGA Secretariat acknowledges this struggle and, over the past two years, has added lightweight access management practices to the DPG Standard as legitimate means of defence, demonstrating how the Standard itself is evolving to address real-world pressures on DPGs without compromising their core commitment to openness:

> "To balance openness with the need to prevent abuse and ensure secure access, we will allow logins, registrations, API keys, and API throttling, provided these measures do not discriminate against any users (including discrimination against any person or group of persons or geographical location or restricting anyone from making use of the content/data in a specific field of endeavour).

However, the problem runs deeper than that. The 'paradox of open' manifests itself because the Open Definition was developed in 2005, in a pre-AI era. It optimises for unconstrained, downstream supply-side open sharing, which—in a landscape of extreme power asymmetry—enables unilateral value extraction rather than equitable data sharing. The current open framework is simply no longer sufficient to govern the structural realities of massive-scale machine consumption. That is also why open data licensing innovations such as the Nwulite Obodo Open Data Licence (NOODL) are so controversial: by introducing sharing-back obligations for data users in high-income countries, the license requires them to share benefits with the data community. It doesn't treat a researcher in Nairobi the same way it treats a publicly listed company in the US, thereby filling an equity gap in data governance. As Dr. Melissa Omino, the license co-creator, contests:

> "[NOODL] rebuilds the relationship between those who give the data and those who take it. That is the kind of complementary mechanism we need: not a retreat from openness, but a more honest version of it, one that protects reciprocity by design rather than asking communities to trust that it will appear later.

However, under the current open definition, NOODL fails to comply.

Purpose and Scope of the Playbook

The intent of this overview is to bring conceptual clarity to a highly contested debate. It does not seek to judge the protective actions taken by DPG product owners and digital commons stewards. Rather, it aims to make visible the struggle product owners face and categorise the defensive mechanisms currently at their disposal. It also includes mechanisms that actively invite reciprocity (marked in pale blue). This overview is by no means complete. Our ongoing conversations with DPG product owners, informed by this playbook and the wider ecosystem debate, will further shape the good practices recommended in the DPG Standard and the guidance provided by the DPG wiki. The above-mentioned and future updates ensure the Standard continues to evolve as a living, responsive framework grounded in the actual defence and reciprocity challenges faced by DPG stewards.

As we document, engaging in a reactive, code-level "cat-and-mouse game" against malicious scrapers by exercising control through client interrogation or data poisoning ultimately punishes human users, undermines data integrity and closes off DPGs. Effective, long-lasting defence of the open ecosystem requires shifting strategies upward and separating legal and technical questions of openness from questions of how resources are governed, how risk is managed, and accountability achieved. Throughout this playbook, each defensive mechanism is assessed not only for its effectiveness but also for its alignment with the DPG Standard. This dual lens ensures that product owners can make informed decisions that protect their projects while preserving their eligibility for recognition as digital public goods.

This is an ecosystem challenge that requires solutions at the ecosystem level. Those shared solutions must move toward the macro-level, legal, and public governance layers. DPGA members are already working towards those solutions by developing sectoral guidance and attribution frameworks under the CC Signals project (Creative Commons) and building alliances around the idea of human knowledge infrastructure as critical digital infrastructure, which requires long-term public funding (Open Knowledge Foundation and Wikimedia Foundation).

Ensuring our shared knowledge infrastructure remains open, human-curated and resilient is the foundation for trust in times of fracturing societies. It's a collective responsibility beyond individual DPG product owners. At the DPGA Secretariat, we're committed to using our platform to collectively change the rules of the game with our members and partners.

Layers of Defence and Reciprocity

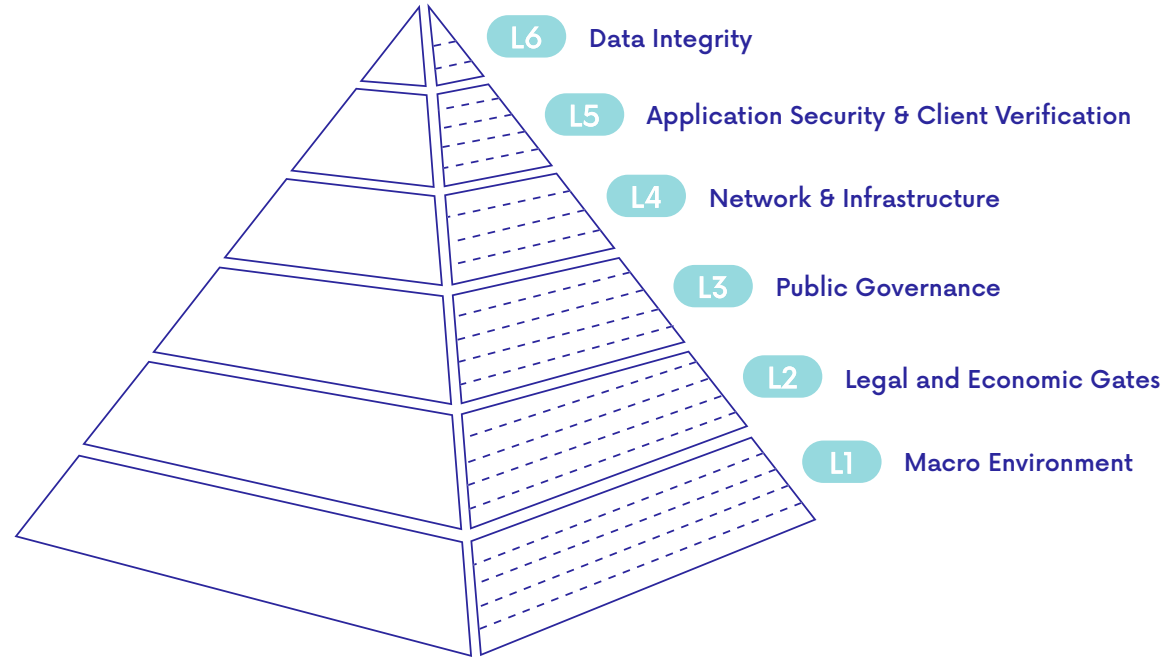

The "defence and reciprocity pyramid" provides an overview of the different levels on which legal, regulatory, governance and technical measures operate.

Layer 1: Macro Environment

Market rules, statutory taxes, and licensing frameworks that operate outside website code.

Layer 2: Legal and Economic Gates

Contracting terms, server-level financial or authentication gates (paywalls, APIs) that control entry.

Layer 3: Public Governance

Non-coercive, machine-readable signs broadcasting usage preferences.

Layer 4: Network & Infrastructure

Systems evaluating how traffic connects to filter out blunt-force bots.

Layer 5: Application Security & Client Verification

Interrogating browser environments and mutating code structure in real-time.

Layer 6: Data Integrity

Structural transformations of the data asset itself, either to poison it or isolate it.

The overview also includes an assessment of whether the mechanism complies with the open definition and, by extension, the DPG Standard since compliance with the Standard is a prerequisite for any digital solution to be formally recognised as a digital public good. Note that defensive mechanisms are usually combined across layers. Some of the mechanisms mentioned, which are classified as a grey area or non-compliant with the openness definition, can also be deployed in parallel with open access pathways, thereby keeping the resource open. An example is **Wikimedia Enterprise**, which introduces a dedicated API for paying clients, providing fast, real-time, and structured access to Wikimedia content; however, the same data is also available for free, though not optimised for high-volume AI scrapers.

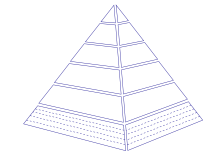

Layer 1: Macro Environment

Category	Mechanism Name	Description	Example	Open Definition and DPG Standard Status
Regulatory	Open data as critical digital infrastructure	Designating human-curated, open data infrastructure as critical infrastructure, which triggers maintenance mandates and public funding via several vehicles (sovereign grants, direct subsidies, etc.)	Policy proposal by the Open Knowledge Foundation and the Wikimedia Foundation	**Acceptable**
Regulatory	Data Excise Taxes & Dividends	Macro-economic sovereign taxes levied against companies harvesting public data, directly redistributed back to citizens (policy proposal)	State-level "Data Excise Taxes" charging a monthly per-user fee to massive tech platforms	**Grey Area**
Legal	Copyleft Data-to-Model Licenses	Legal frameworks dictating that if a dataset is used to train an AI, the resulting model parameters must be open-sourced.	Releasing a text corpus under GNU GPL-style clauses, e.g. Open Data Commons Open Database License (ODbL) used by OpenStreetMap or licensing innovations such as the Contextual Copyleft AI License (CCAI)	**Grey Area**
Legal	License Innovation	Equitable, community-centred, tiered open data licenses that apply context-sensitive terms, geographic differentiations, and benefit-sharing obligations.	Nwulite Obodo Open Data Licence (NOODL)	**Fails**
Legal	Extended Collective Licensing (ECL)	Unified clearinghouses that license vast institutional data blocks to AI developers, ensuring automated royalty streams.	The Copyright Clearance Centre's (CCC) AI Systems Training License	**Fails** **Acceptable if used in parallel**
Legal	Closed license	Licenses that exert control over how content can be used.	CC BY-NC-ND (Creative Commons Attribution Non-Commercial Non-Derivatives License)	**Fails**

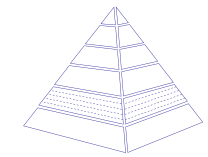

Layer 2: Legal and Economic Gates

Category	Mechanism Name	Description	Example	Open Definition and DPG Standard Status
Legal	API Access Deals	Commercial agreements offering paid high-volume data access while keeping the raw dataset free elsewhere.	Wikimedia Enterprise	**Fails** **Acceptable if used in parallel**
Legal	Resource Bartering (Inference Credit Swaps)	An organisational contract trading raw data directly for high-tier AI inference tokens or corporate compute power.	Major closed content holders, such as Reddit, use versions of this	**Fails** **Acceptable if used in parallel**
Legal	Terms of Service (ToS) Clauses	Contractual rules governing website usage that explicitly prohibit automated data harvesting and scraping.	Mozilla Common Voice Dataset	**Grey Area**
Technical	Data Flywheel API Locks	A structural loop where high-volume data access is conditioned on the developer feeding model metadata back to the host.	Forcing an autonomous vehicle company to feed edge-case labels back to a mapping database	**Fails** **Acceptable if used in parallel**
Legal	Paywalls	Moving public datasets behind a mandatory financial barrier or paid user authentication tier.	Placing the raw dataset download behind a $5/month premium subscription	**Fails**
Technical	Request-to-access (Gated Data)	A data access model requiring users to submit a formal application detailing identity and intent before access.	Rights Tracker and Mozilla Data Collaborative use data access forms	**Fails (if intent is used to deny access)**

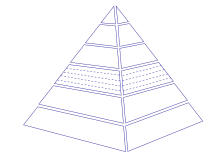

Layer 3: Public Governance

Category	Mechanism Name	Description	Example	Open Definition and DPG Standard Status
Governance	Robot Policy (robots.txt)	A standardized text file placed on a website to instruct web robots which pages they are permitted to index.	robots.txt disallowing specific user-agents	**Acceptable**
Governance	Preference Signals	Machine-readable protocols such as Universal Opt-Out Mechanisms (UOOMs) or automated rights-reservation embedded in web content that broadcast a publisher's explicit data-use restrictions.	Global Privacy Control (GPC) or the W3C Community Group's tdm-reservation protocol	**Acceptable**
Governance	Attribution Standards	Forcing a behavioural loop where AI outputs must link back to and cite the original source datasets that informed them.	Coalition for Content Provenance and Authenticity (C2PA) standards or web search AI engines embedding mandatory source links	**Acceptable**

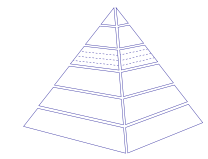

Layer 4: Network & Infrastructure

Category	Mechanism Name	Description	Example	Open Definition and DPG Standard Status
Technical	Web Application Firewall (WAF)	A network security solution that monitors and filters traffic to block known automated scraper bots.	AWS WAF bot control rules	**Grey Area**
Technical	Bot-detection heuristics	Algorithmic analysis of connection attributes and request patterns to mitigate automated traffic in real-time.	Flagging a client that requests 500 pages per second with perfectly uniform timing	**Grey Area**
Technical	Bot Blocking	Technical server configurations used to explicitly deny access to specific known AI crawler user-agents.	Blocking the GPTBot or ClaudeBot user-agents in server configurations	**Grey Area**
Technical	API Rate Limiting	Restricting the number of requests an IP address can make within a timeframe to prevent mass data harvesting.	Limiting unauthenticated API requests that significantly exceed average traffic	**Grey Area**

Technical	API Keys	Providing authentication mechanisms to track and control API usage and revoke access if abused.	User API keys that expire, have fixed request paramenters	**Grey Area**
Technical	Regional Traffic Prioritisation	Restricting web traffic based on geographic origin to prioritise local users over foreign server farms.	Throttling or blocking non-domestic IP ranges during peak local hours	**Grey Area**

Layer 5: Application Security & Client Verification

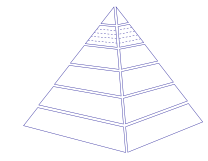

Category	Mechanism Name	Description	Example	Open Definition and DPG Standard Status
Technical	In-Browser Client Verification	Non-intrusive background browser challenges are used to verify that a client is a human-operated program.	Cloudflare Turnstile	**Grey Area**
Technical	CAPTCHAs	Interactive challenge-response tests used to determine whether a user is human to block automated scripts.	Google reCAPTCHA v3	**Grey Area**
Technical	Proof-of-work challenges	Cryptographic puzzles a browser must solve before accessing a page, raising the computational cost for scrapers.	Friendly Captcha	**Grey Area**
Technical	Dynamic HTML / DOM Shuffling	Constantly scrambling the underlying website source on each deployment build to break automated data parsing scripts.	Meta/Facebook dynamically changing CSS class names on every page load	**Fails**

Layer 6: Data Integrity

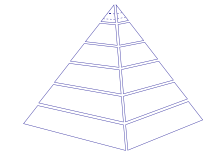

Category	Mechanism Name	Description	Example	Open Definition and DPG Standard Status
Technical	Data Poisoning / Perturbations	Intentionally altering dataset files with invisible artefacts to corrupt or break downstream AI training models.	Using tools like Glaze or Nightshade on images or text	**Fails**
Technical	Canary Data / Honeytokens	Injecting unique, fabricated data into a dataset to uniquely track and prove unauthorised downstream scraping.	Seeding a text dataset with fake names or fictitious dictionary definitions	**Fails**

Compliance Overview

As defensive mechanisms move from macro-governance (layers 1-3) down into infrastructural solutions and data alterations (layers 4-6), compliance with the Open Definition and the DPG Standard declines. The most aggressive technical protections fundamentally violate the core tenets of open data. Positive reciprocity mechanisms are only found in layers 1-3. This means the most promising mechanisms that protect openness while defending the resource are systemic in nature and operate at the macro level rather than the code level.

	Acceptable	Grey Area	Fails
Layer 1: Macro Environment	● ●	● ●	● ●
Layer 2: Legal and Economic Gates	● ● ●	●	● ●
Layer 3: Public Governance	● ● ●		
Layer 4: Network & Infrastructure		● ● ● ● ● ●	
Layer 5: Application Security & Client Verification		● ● ●	●
Layer 6: Data Integrity			● ●

The DPG Defence Playbook

Navigating the Paradox of Open to Mitigate AI Exploitation

If you think there's a mechanism that is missing from this overview or a classification is inaccurate, please reach out to **hello@digitalpublicgoods.net**. Your feedback will help inform future updates to the DPG Standard and its accompanying guidance.

digitalpublicgoods.net