DELIVERABLE 2.4

Saufex

Institutional Recommendation Report

Project acronym:	*SAUFEX*
Due date of deliverable:	*Month 30*
Actual submission date:	*2 Oct, 2026*
Work package:	*WP2*
Lead beneficiary:	*DEB*
Dissemination level:	*PU*
Status:	*Issued*

Funded by
the European Union

Document history

Revision	Date	Description
1	1 Sept, 2026	First issue
2	14 Sept, 2026	Partners feedback
3	2 Oct, 2026	Final issue

Table of Contents

Abbreviations and acronyms 1

Executive Summary 6

1. Introduction 7

2. Methodology 9
2.1 Why a single-actor case study 10
2.2 A note on terminology 10
2.3 The evidence base 11
2.4 Cross-report entity resolution and event deduplication 12
2.5 AI-assisted data collection and codification 13
2.6 Safeguarding GDPR rules 14
2.7 How sources are assessed: a two-axis provenance model 14
2.8 The analytical taxonomy: Information Manipulation Sets 15
2.9 Three IMSs, one Threat Actor 16
2.10 Cleaning the actor picture before counting 18
2.11 How the framework is scored against the record 20
2.12 Two counts of response, deliberately distinguished 23
2.13 Scope and limitations 24

I. The threat actor and its footprint 24

3. The threat actor in brief: SDA case analysis 24
3.1 SDA and the privatisation of Russian information operations 24
3.2 The threat actor at a glance 28
3.3 Three Information Manipulation Sets, one Threat Actor 29
3.4 How SDA operates 30
3.5 What the response has achieved, in brief 31

4. How SDA came to be documented 31
4.1 The existing literature 31
4.2 Internal leaks and extracted data 33

5. Operations: process and record 33
5.1 How an SDA operation runs 34
5.2 What was executed and what was only planned

36
II. The response framework 38

6. Rationale for a multi-report approach **38**

7. The ecosystem inventory: targets, intermediaries and responders **40**

8. The reporting matrix 42
8.1 The matrix 42
8.2 The bulletproof-hosting problem 43
8.3 Viable options and intervention toolkits 44
8.4 Service withdrawal on contractual grounds 45
8.5 Standardisation of the reporting pipeline 46
8.6 The trusted-flagger model 47

9. Evidence standards for court-admissible action 47

III. Testing the framework against the SDA case 49

10. The response record and the scale mismatch **49**

11. Findings: successes and limitations **54**

12. Countering the hardest element: sanctioning bulletproof hosting (the AEZA case) 56

IV. Lessons learned and recommendations 58

13. Lessons learned 58
13.1 Knowledge was not the constraint 58
13.2 The EU response barely reached the middle of the pyramid 59
13.3 Each measure was a single strike 60
13.4 What worked was decentralised 61
13.5 State instruments were not carried through 61
13.6 A contractor can be hurt where a business is hurt 61

14. Recommendations 62
14.1 The approach 62
14.2 Designate the infrastructure, in step with commercial action 63
14.3 Stop payments from designated actors 64

14.4 Notify brand owners and providers 64
14.5 Fund and mandate monitoring after each measure 65
14.6 Share measures and outcomes under one operation identifier 65
14.7 Counter-respond at national level 65
14.8 Adopt a common evidence standard 66

14.9 Create a notice route against LLM grooming 66
14.10 Safeguards 66

15. Conclusion 67

Figures and Tables 68

Annexes 69
Annex A - The list of the 88 sources 69
Annex B - The 33 binding measures 76

Abbreviations and acronyms

Abbreviation	Meaning	Description
ABCDE	**Actor, Behaviour, Content, Degree, Effect**	An analytical framework referring to Actor, Behaviour, Content, Degree, Effect, used to structure FIMI analysis.
AFNIC	**Association Française pour le Nommage Internet en Coopération**	The French country-code registry, which operates its own domain-dispute procedure.
AI	**Artificial Intelligence**	Systems or models capable of performing tasks that typically require human intelligence, including pattern recognition, content generation, classification, and predictive analysis.
ANO	**Autonomous Non-Profit Organisation**	A Russian legal form; used here in the name ANO Dialog.
ccTLD	**Country-code top-level domain**	A top-level domain assigned to a country or territory, such as .fr or .re, whose registry sets its own dispute procedure.
CERT	**Computer Emergency Response Team**	A national or sectoral body responsible for responding to cyber incidents.
CIB	**Coordinated Inauthentic Behaviour**	Meta's term for coordinated activity by accounts that misrepresent who is behind them.
CMS	**Content Management System**	Software used to publish and manage a website, whose user accounts can link assets to one another.
CTI	**Cyber Threat Intelligence**	The discipline and data standards from which the Information Manipulation Set construct is adapted.
DAD-CDM	**Defense Against Deception – Common Data Model**	The DAD-CDM open project focuses on countering disinformation and related threats, including misinformation, influence operations, and foreign information manipulation. Its mission is to strengthen defence and resilience by developing standards and practical guidance that support both real-time response and long-term risk management.

Abbreviation	Meaning	Description
DFRLab	**Digital Forensic Research Lab**	Research body of the Atlantic Council, and a source in the corpus.
DISARM	**Disinformation Analysis and Risk Management Framework**	An open-source framework designed for describing and understanding the behavioural parts of disinformation and FIMI.
DOJ	**United States Department of Justice**	Authority behind the September 2024 domain seizures.
DSA	**Digital Services Act**	Regulation (EU) 2022/2065 on a single market for digital services.
DSC	**Digital Services Coordinator**	The national authority designated under the Digital Services Act.
EC	**European Commission**	The executive branch of the European Union responsible for proposing legislation, implementing decisions, and upholding EU treaties.
EEAS	**European External Action Service**	The European Union's diplomatic service supporting the High Representative in conducting the EU's foreign and security policy.
EU	**European Union**	
FBI	**Federal Bureau of Investigation**	United States authority, author of the affidavit supporting the September 2024 seizures.
FCDO	**Foreign, Commonwealth and Development Office**	United Kingdom department responsible for its sanctions designations.
FIMI	**Foreign Information Manipulation and Interference**	Hostile information activities, often state-sponsored, aimed at undermining democratic processes and institutions.
GDPR	**General Data Protection Regulation**	Regulation (EU) 2016/679, governing the processing of personal data in the codified material.

Abbreviation	Meaning	Description
IMS	**Information Manipulation Set**	A collection of adversarial behaviours, tools and TTPs presumed linked to the same Threat Actor or group of Threat Actors, which may be unknown; concept defined by VIGINUM (2026), interoperable with STIX 2.1, adopted by the EEAS, the G7 Rapid Response Mechanism and the January 2026 FIMI working group.
IOC	**Indicator of Compromise**	A technical artefact, such as a domain or an address, evidencing an operation.
IP	**Internet Protocol**	The addressing system of the internet; IPv4 and IPv6 are its two versions, and blocks of addresses are announced as prefixes.
ISAC	**Information Sharing and Analysis Centre**	Sectoral body for exchanging threat information; FIMI-ISAC is the one for information manipulation.
ISD	**Institute for Strategic Dialogue**	Research body and a source in the corpus.
ISO	**International Organization for Standardization**	An independent international body that develops and publishes global standards to ensure quality, safety, efficiency, and interoperability across sectors.
KPI	**Key Performance Indicator**	Used here for the reach figures the operation reported to its sponsor.
LLM	**Large Language Model**	A generative model that absorbs published material and reproduces it in answers.
OASIS	**Organization for the Advancement of Structured Information Standards**	An international standards consortium that develops open standards for information exchange, including STIX and related cybersecurity frameworks.
OCCRP	**Organized Crime and Corruption Reporting Project**	Investigative consortium and a source in the corpus.
OFAC	**Office of Foreign Assets Control**	The sanctions authority of the United States Treasury.
OFSI	**Office of Financial Sanctions Implementation**	The sanctions authority of HM Treasury, United Kingdom.

Abbreviation	Meaning	Description
OpenCTI	**Open Cyber Threat Intelligence**	Open-source platform for managing and analysing cyber and disinformation threat intelligence.
OSINT	**Open-source Intelligence**	Intelligence derived from publicly available sources, including media, social platforms, public records, and online databases.
PFKI	**Presidential Foundation for Cultural Initiatives**	Russian state fund created in 2021, cited as a source of grants to a PR agency.
RFE/RL	**Radio Free Europe / Radio Liberty**	Broadcaster and a source in the corpus.
RRN	**Reliable Recent News**	One of SDA's three Information Manipulation Sets, renamed twice and now Researchers and Reporters Network.
SAUFEX	**Secure Automated Unified Framework for EXchange**	The Horizon Europe project under which this deliverable is produced.
SBU	**Security Service of Ukraine**	Ukrainian authority notified in one of the recorded referrals.
SDA	**Social Design Agency**	The Russian company that is the subject of the case study.
SDO	**STIX Domain Object**	A primary object type within STIX used to represent high-level threat intelligence concepts such as threat actors, campaigns, or attack patterns.
STIX	**Structured Threat Information Expression**	A standardized language and data model for describing, exchanging, and storing cyber and hybrid threat intelligence.
STIX 2.1	**Structured Threat Information Expression (version 2.1)**	Standardised data format for representing and sharing threat intelligence.
TLD	**Top-level domain**	The final segment of a domain name, such as .com or .cc.

Abbreviation	Meaning	Description
TLP	**Traffic Light Protocol**	A classification framework used to indicate how sensitive information may be shared, based on colour-coded dissemination restrictions (for example TLP:RED, AMBER, GREEN, CLEAR).
TTP	**Tactics, Techniques, and Procedures**	The behavioural methods and operational patterns used by threat actors to plan, execute, and sustain malicious or manipulative activities.
UDRP	**Uniform Domain-Name Dispute-Resolution Policy**	An ICANN policy under which a trademark owner can obtain transfer of a domain confusingly similar to its mark, decided in administrative proceedings rather than in court.
URL	**Uniform Resource Locator**	A standardised web address used to identify and access specific resources on the internet.
VIGINUM	**Service de vigilance et de protection contre les ingérences numériques étrangères**	The French state service against foreign digital interference, and the source of the Information Manipulation Set construct.
WIPO	**World Intellectual Property Organization**	United Nations agency whose Arbitration and Mediation Center is the principal accredited provider administering UDRP proceedings.
WP	**Work Package**	A defined unit of work within a project structure, outlining specific objectives, tasks, deliverables, timelines, and responsibilities.
WP2	**Work Package 2**	Standardisation, Development of Automated Codification Processes, and Training.
WP4	**Work Package 4**	Development, Testing, and Validation of AI Tools.

Executive Summary

D2.4 delivers recommendations for institutions on how to respond to foreign information manipulation and interference (FIMI). It contributes most directly to the project's second objective, to de-centralise and democratise processes around FIMI analysis and response.

The deliverable proposes a response framework that identifies seven elements of an operation against which action can be taken, and the parties able to act on each. It tests the framework against the public record of one threat actor, the Russian Social Design Agency (SDA) and its partner companies Structura National Technologies and ANO Dialog. The record was assembled from 88 sources into a single registry of entities, infrastructure and response measures, in which each enforcement act is counted once, however often it was reported.

Main findings

- **Knowledge was not the constraint.** Four governments attributed the operation, three jurisdictions designated it and a US court authorised the seizure of its domains. The actor still operates.
- **The EU response barely reached the middle of the operation.** Using David Bianco's Pyramid of Pain, adapted from cyber threat intelligence, the report ranks each part of the operation by how costly it is for the actor to replace. Designations reached the organisation, and takedowns reached domains and accounts that were replaced within a day. Hosting, advertising and payment, which the operation cannot replace overnight, drew 13 recorded binding measures between them, every one from a commercial provider or from a US or UK designation. The Union's one listing of a provider the operation had used, Stark Industries in May 2025, falls outside the record, and Aeza, which still carries the operation, has not been listed.
- **Each measure was a single strike.** Of 33 binding measures, 11 have no recorded outcome, and 5 of the 22 with an outcome led the operation to adapt and continue. The record rarely shows anyone watching what the operation rebuilt, and no binding measure was aimed at an operation as such.
- **What worked was decentralised.** Brand owners recovered all four domains they disputed. Between July and November 2024, hosting, transit, cloaking, registrar and telecommunications providers withdrew service under their own terms after two research organisations approached them, with the strongest effects in the record.

- **State instruments were not carried through.** The advertising platform kept receiving payments from the operation after the EU and US designations. Aeza, the operation's bulletproof host, was designated by the United States and the United Kingdom, yet reaches the internet mainly through a German transit provider, in the one jurisdiction that has not designated it.

Recommendations

The report asks institutions to make the routes that worked systematic, and to do what only institutions can do. Its eight recommendations follow a loop of respond, monitor and counter-respond, and each names the body that would act: the Council and the Commission on designating infrastructure; sanctions authorities, the Commission and Resilience Councils on payments to designated actors; Resilience Councils and Digital Services Coordinators on notifying brand owners and providers; the Commission and Member States on funding monitoring; the EEAS, FIMI-ISAC and the G7 Rapid Response Mechanism on a shared operation identifier; national CERTs and Digital Services Coordinators on the national follow-up; Digital Services Coordinators and the European Board for Digital Services on an evidence standard; and the AI Office and the Commission on LLM grooming. A section on safeguards sets out the conditions that keep these measures proportionate.

The evidence measures effects on infrastructure. No source measures effects on audiences, and a third of the binding measures have no recorded outcome. A single case shows where a response fails, and cannot show how often. The next step is to apply the same test to a second actor, such as Storm-1516.

1. Introduction

This deliverable is produced under Work Package 2 of the SAUFEX project, whose objectives concern standardisation, the development of automated codification processes, and training. D2.4 is the work package's institutional output: where earlier deliverables addressed how foreign information manipulation is detected and described, this one addresses what is done about it once it has been.

D2.4 contributes most directly to the second SAUFEX objective: to de-centralise and democratise processes around FIMI analysis and response. The Description of Action frames that objective around three points: the difficulties Member States face in detecting, attributing and sanctioning FIMI; the weight that responses carry for fundamental rights; and two mechanisms to address both, Resilience Councils acting between state and non-state actors, and national Digital Services Coordinators equipped to handle evidence, notice and takedown without resorting to censorship. The findings of this deliverable bear directly on that objective. The responses with the strongest recorded effects against the Social Design Agency came from brand owners and from commercial providers acting at the prompting of researchers, and none of them required a court or a regulator. Part IV draws the institutional consequence: institutions should make those decentralised routes systematic and use the instruments that only they hold. Its recommendations are addressed to Resilience Councils and Digital Services Coordinators alongside EU institutions, national authorities and intermediaries, and draw on the Resilience Council model set out in D1.1.

That division reflects an imbalance in the field. A review of thirty-six response and countermeasure frameworks, conducted as background to this deliverable, found fourteen in operational use, sixteen at the stage of doctrine or early adoption, and six existing only on paper. The analytical half of the discipline is mature and machine-readable, with established standards for describing threats and the behaviours behind them. The response half consists largely of menus of options, but almost no framework separates the effect of a response from the effect of the campaign it answers. Detection has been solved to a far greater degree than response.

D2.4 addresses that gap in two movements. It proposes an institutional response framework: a model of who must report to whom, in what form, and against which parts of an operation. It then tests that framework against the complete published record of a single threat actor, the Social Design Agency, reconstructed from 88 sources into a consolidated registry of entities, relationships, infrastructure and

response measures. The purpose of the test is not to describe the actor, which has been done many times over, but to establish where a response framework holds and where it fails when it meets an operation that replaces its assets within a day and routes through providers whose business is to ignore complaints.

The timing is deliberate. During 2026 the analytical vocabulary of the field consolidated around the Information Manipulation Set, defined by VIGINUM in January, applied in the common operational picture published by a working group of European and North American organisations in the same month, adopted by the European External Action Service in its fourth threat report in March, and endorsed by the G7 Rapid Response Mechanism in May. That consolidation gives the response side a shared unit of analysis for the first time. This deliverable adopts it, and asks what follows for practice.

The document is organised in four parts. Part I sets out the threat actor, how it came to be documented, and how its operations run, establishing the footprint against which any response must be measured. Part II sets out the response framework itself: the case for reporting to several recipients at once, the inventory of parties who can act, the matrix of what each receives, and the evidentiary standard that reporting must meet. Part III tests the framework against the record, examining what the response consisted of, what it achieved, and the case of the hosting provider that has proved hardest to reach. Part IV draws the lessons and sets out eight recommendations, each addressed to the body that would have to act. A Methodology section precedes Part I and should be read first, since the counting rules it establishes govern every figure that follows.

Two limits are stated at the outset. This is a single case, and a single case can show that a response framework fails at particular points; it cannot establish how often those failures occur in general. And the evidence base measures effectiveness against infrastructure — domains seized, accounts suspended, services withdrawn — because no source in the corpus measures the effect of any measure on any audience. What this deliverable can say about influence is therefore bounded, and that boundary is itself among its findings.

2. Methodology

The deliverable 2.4 Institutional recommendation is not intended as a survey of FIMI countermeasures in the abstract, but rather a stress test from which conclusions and recommendations can be drawn. D2.4 proposes an institutional response framework, a model of who must report to whom and how, and identifies which elements of a FIMI operation's kill chain[1] should be targeted in order to effectively contain the operation rather than merely document it. D2.4 takes that framework off the page and runs it against one real, fully-attributed actor to see where it holds and where it breaks.

In D2.4 we measure the response framework against the malign activities of the Social Design Agency (hereinafter SDA), a Russian IT and PR company, and ask a single empirical question: did the response reach the whole strategy of SDA, or only a fraction of it? We attempt to answer it from an 88-source evidence base. The core claim starts from observing that detection and attribution are saturated, but response was partially ineffective.

2.1 Why a single-actor case study

A framework that lists "report to the registrar, report to the host, report to the platform" might look complete on paper. But its weaknesses become visible when a real operation is laid against it. For instance, when an actor spans dozens of jurisdictions, re-mints its infrastructure within a day, and routes through hosts whose business model is ignoring abuse reports.

SDA was chosen as that actor because it is, paradoxically, one of the most exposed FIMI actors in existence: despite two major leaks (September 2024[2]; May 2026[3]), the reports and attributions of four governments, France (VIGINUM), the United States (DOJ affidavit and OFAC designations),[4] the United Kingdom (FCDO/OFSI sanctions),[5]

[1] In general terms, a kill chain can be described as a systematic process used to identify, target, and neutralize a threat.
[2] "Leaked Records Detail Vast Russian Influence Campaign Targeting Ukraine, EU", RFE/RL (Schemes), with Süddeutsche Zeitung, NDR, WDR and VSquare, 17 September 2024, https://www.rferl.org/a/russia-ukraine-disinformation-propaganda-campaign/33123498.html.
[3] M. Laine, M. Vunš, I. Lozovsky, "Leaked Documents Reveal Russian 'Cognitive Strikes' Against the West — Including Islamophobic 'Pig Head' Attacks in Paris", OCCRP and Delfi, 24 May 2026, https://www.occrp.org/en/investigation/leaked-documents-reveal-russian-cognitive-strikes-against-the-west-including-islamophobic-pig-head-attacks-in-paris.
[4] Affidavit in Support of Seizure Warrant, United States v. Certain Domains, Case No. 24-mj-1395, United States District Court for the Eastern District of Pennsylvania, 4 September 2024, https://www.justice.gov/d9/2024-09/doppelganger_affidavit_9.4.24.pdf.
[5] Foreign, Commonwealth & Development Office, "UK sanctions Putin's interference actors",

and Germany (Federal Foreign Office technical analysis),[6] and dozens of independent investigations have mapped it, it still operates.

2.2 A note on terminology

To avoid confusion, six specific terms are used in one sense only throughout this report. Four of them describe the data and two describe the analytical framework. The corpus is the evidence base of 88 publicly available sources. From each source, information was extracted through a codification, one structured file per report. Those codifications were then consolidated and resolved into a single registry, which is what every figure in this deliverable is counted from. These three terms form a chain of refinement from raw data to the codified registry.

The fourth term, dataset, refers only to third-party files of indicators of compromise published by other organisations. The distinction matters because a figure drawn from a dataset cannot be compared with one drawn from the registry. The remaining two terms, level and element, belong to the analytical framework and are defined in Table 1.

Table 1. The controlled vocabulary of this deliverable.

Term	What is describes
corpus	the 88 source reports
codification	the per-report structured extraction (one file per report)
registry	the consolidated, resolved collection (entities, relationships, decisions, observables) assembled from the codifications
dataset	only third-party published IOC files (VIGINUM CSVs, Meta IOC sets), each with its own provenance
level	one of the six levels of the IMS schematisation, from political to technical
element	one of the seven parts of an operation against which a report can be made

28 October 2024; HM Treasury, Office of Financial Sanctions Implementation, designation of the Social Design Agency, unique ID RUS2292, Russia (Sanctions) (EU Exit) Regulations 2019.

[6] German Federal Foreign Office, Technical Report on the "Doppelgänger" Disinformation Campaign — Analysis of German-language activities, 5 June 2024, https://www.auswaertiges-amt.de/resource/blob/2682484/2da31936d1cbeb9faec49df74d8bbe2e/technischer-bericht-desinformationskampagne-doppelgaenger-1--data.pdf..

2.3 The evidence base

The case is built on 88 different sources that fall into five types, each contributing a different kind of proof: information and details from the leaked files of SDA reported by media outlets, the government documents and decisions, the platforms' detections and enforcement, the technical data analyzed by FIMI and OSINT organizations, and the investigations published by online media. One further document had been codified twice under two identifiers and is counted once.

The two major leaks on SDA do not constitute per se two of the 88 sources. Information deriving from leaked material was not directly included in the evidence base as raw documents, but rather filtered through the multiple publications that quoted or mentioned such material, providing useful corroborating evidence. The 2024 "Kremlin Leaks", obtained by Delfi Estonia and distributed in February, enter the corpus through Meduza and France 24.[7] The Social Design Agency's own internal files, reported in September 2024, enter it through VSquare with Delfi Estonia and Frontstory, and through RFE/RL.[8] The second tranche of SDA leaks, also obtained by Delfi Estonia, was shared in 2026 with OCCRP and VSquare.[9] In addition, the Gnida Project, Süddeutsche Zeitung, Bloomberg and Forbidden Stories, among others, quoted the leaked material too.[10] Table 2 sets out the types of source in the corpus and what each establishes.

Table 2. Types of source in the corpus and what each establishes.

Type of source	What it establishes
The leaks (actor's own files)	Organisation chart, tasking chain, budgets, project plans, self-reported reach
Government action	DOJ affidavit, OFAC/OFSI/EU sanctions, indictments and convictions

[7] Meduza, "'Latching onto successful projects': Leaked documents suggest Kremlin spin doctors are presenting popular movies and TV to Putin as propaganda wins", 27 February 2024; France 24, "'Kremlin Leaks': Files detail Putin's €1 billion propaganda effort ahead of presidential vote", 27 February 2024.
[8] VSquare, "Kremlin Leaks: Secret Files Reveal How Putin Pre-Rigged his Reelection", February 2024; RFE/RL (Schemes), "Leaked Records Detail Vast Russian Influence Campaign Targeting Ukraine, EU", 17 September 2024;
[9] OCCRP and VSquare, "Leaked Documents Reveal Russian 'Cognitive Strikes' Against the West - Including Islamophobic 'Pig Head' Attacks in Paris", 24 May 2026.
[10] Gnida Project, "Recent Social Design Agency leak points to Storm-1516 coordination with the Presidential Administration", 30 May 2026; Bloomberg, "Russia's Wiki Warfare Tries to Distort Reality, Documents Show", 22 June 2026; Süddeutsche Zeitung, "So flutet Russland das Netz mit Desinformation", 16 September 2024; Forbidden Stories, "Propaganda Machine: Secret documents reveal Russia's foreign influence strategy across three continents", 14 February 2026.

Type of source	What it establishes
Platform enforcement	Meta CIB takedowns and indicator-of-compromise sets
Technical / OSINT analysis	Reports by VIGINUM, Qurium, Recorded Future, DFRLab, Graphika
Investigative journalism	Investigations by OCCRP, VSquare, Bloomberg, DW, BIRN, RFE/RL

2.4 Cross-report entity resolution and event deduplication

The same organisation and the same enforcement action may be described in different reports, each in its own words. If the codifications of each report and source were simply added together, every retelling would be counted as a separate item and the totals would be grossly inflated. The codifications were therefore consolidated into a single registry before any figure in this report was stated.

The clearest illustration of this concept is the seizure of Doppelgänger domains by the United States Department of Justice in September 2024.[11] This is one action, but the corpus describes it fourteen times over. Those fourteen accounts name the acting authority nine different ways (variously the DOJ, the FBI, the US government, US authorities, and so on) and provide five different dates for the enforcement. Without an event deduplication, this single seizure of domains would enter the record as nine separate enforcement actions.

Therefore, the consolidation proceeds in three steps: the entity resolution merging all variants and names for the same organisation, person or asset, the merge of the same relationship reported by different sources, and the deduplication of the same enforcement actions and responses recorded by multiple sources. However, the process of consolidation and entity resolution was not left entirely to the AI: merges and deduplications that were uncertain and would open a question were reviewed by the analysts for the final judgement.

A second pass was then made over the resulting decisions, because the deduplication had not merged every retelling of the same act. The EU listing of July 2023 remained as six records, the US designation of March 2024 as seven and the September 2024 seizure as seven. Each record was therefore mapped by hand to the real-world act it describes, and records that on review concerned other actors, or no binding measure at all, were removed. Records the registry had flagged as not concerning SDA were

[11] Affidavit in Support of Seizure Warrant, United States v. Certain Domains, Case No. 24-mj-1395, United States District Court for the Eastern District of Pennsylvania, 4 September 2024.

also read, which recovered measures the first pass had missed, among them Google's delisting of 47 Doppelgänger domains. The continuous enforcement of a platform, observed by several researchers, is counted as one standing act. The unit of counting is the act, which is counted when it addresses an element of SDA's operation, whether or not it names SDA. The pass leaves 33 binding measures. Where the records of one act disagree on its outcome, the assessment with the most corroborating reports is kept and the act is marked as contested. For eight acts the outcome was coded by the analysts rather than taken from the registry: seven withdrawals of service in 2024, coded from what each source records, and the 2025 US designation of Aeza, whose recorded outcome is contradicted by the routing measurement in Section 12. The mapping, the reasons and the rules are supplied with the registry.

2.5 AI-assisted data collection and codification

The collection, extraction, and structuring of data from the 88-source corpus was performed with the support of an AI system (Claude Code by Anthropic), operating under continuous human supervision. The AI was used in codification of four defined tasks: gathering open-source information from public reports, media investigations, and online repositories; extracting entities (organisations, individuals, domains, operations) into a structured codification, one file per report; de-duplicating aliases and cross-referencing claims across sources; and flagging inconsistencies between sources for analyst review. All analytical judgements, such as source reliability ratings, attribution decisions, the separation of executed from planned operations, and every figure reported in this deliverable, were made or validated by the research analyst. The AI functioned as a processing instrument, not as an autonomous analytical agent, in line with the human-in-the-loop principle that this report itself recommends for institutional adoption.

2.6 Safeguarding GDPR rules

Some operational measures are implemented throughout the research process. First, the unit of analysis is the organisation, not the individual. The report analyses the SDA as an operational structure: its entities, infrastructure, campaigns, budgets, and methods, rather than profiling the persons within it. This organisation-level focus minimises the processing of personal data and aligns with the proportionality logic applied in Horizon Europe ethics review: pattern-level analysis of an influence operation serves a strong public interest with minimal intrusion into individual privacy.

In addition, no identification is made beyond the established public record and this report names no individual who has not already been identified in reputable public reporting (government attributions, sanctions listings, indictments, or major investigative publications). It does not combine leaked data with public information in

ways that would enable the re-identification of individuals appearing only in the leaks. Leaked material was used at the processing stage, for the AI-assisted structuring of the actor picture and for the analysts' general understanding of the operation, but personal data or private information appearing exclusively in the leaks is not reproduced, quoted, or disclosed in this deliverable.

Finally, the AI-assisted collection itself was confined to publicly accessible material and conducted in accordance with the transparency and platform-accountability framework of the Digital Services Act. No platform data requiring vetted-researcher access under Article 40 DSA was used, and no scraping in breach of platform terms of service was performed.

2.7 How sources are assessed: a two-axis provenance model

Influence-operation evidence can be misleading because the richest source, the actor's own files, is also the most self-serving. Another example of this comes from another leak regarding the so-called "Company", disclosed by Forbidden Stories and partners.[12] To avoid laundering the actor's marketing into the report's findings, every item is rated on two independent axes:

- **Artefact reliability:** how trustworthy the document or capture is as a record (for instance, a verified court exhibit ranks far above a screenshot from a chat).
- **Claim credibility:** whether the assertion inside is independently corroborated (a leaked plan proves intent, not execution; a self-reported '86M views' figure proves a KPI was claimed, not that it was achieved).

Keeping the axes separate allows us to state that an operation was genuinely planned and budgeted while flagging at the same time that there is no evidence it was ever executed. This is a distinction the report enforces throughout by separating live/executed operations from those only planned or proposed.

2.8 The analytical taxonomy: Information Manipulation Sets

Before any figure is reported, this deliverable needs a vocabulary that keeps distinct what the FIMI community has habitually merged. The malicious digital activities attributed to the Social Design Agency are referred to in the literature, interchangeably, as an "operation", a "campaign", a "network", or even a "Threat Actor". VIGINUM identified this taxonomy uncertainty as a structural obstacle to threat

[12] L. Peruchon, E. Layan, S. Álvarez Jurado, "Propaganda Machine: Secret documents reveal Russia's foreign influence strategy across three continents", Forbidden Stories, 14 February 2026, https://forbiddenstories.org/propaganda-machine-secret-documents-reveal-russias-foreign-influence-strategy-across-three-continents/.

analysis, citing precisely the Doppelgänger case as its lead example.[13] To resolve it, this report adopts the concept of Information Manipulation Set (IMS), defined by VIGINUM as "a collection of adversarial behaviours, tools, and Tactics, Techniques, and Procedures (TTPs) presumed to be linked to the same Threat Actor or group of Threat Actors, which may be unknown". An IMS is not the Threat Actor. It is the tangible manifestation of the actor's activity, sitting at the infrastructural level below the Threat Actor, qualified by two cumulative criteria: **clandestinity** (demonstrable efforts by the operators to prevent the linking of assets to each other and to the actor) and **coordination** (strong technical or behavioural links among the assets, which are presumed to be operated by the same actor).

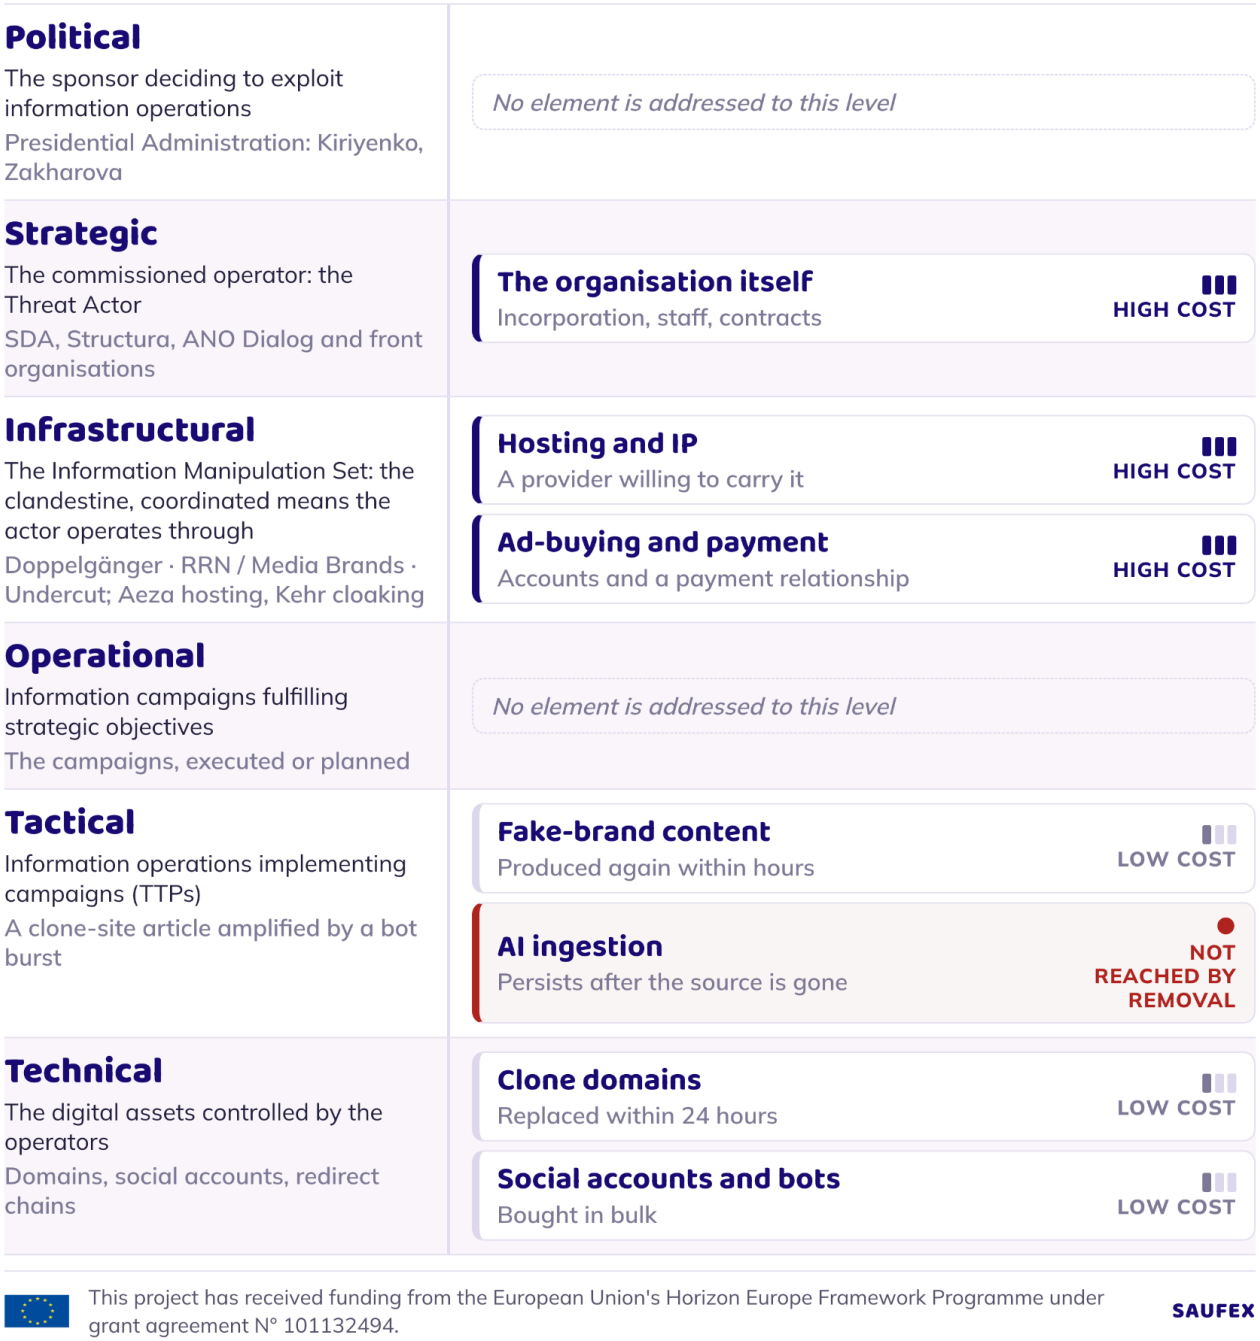
The concept of IMS therefore sits within a schematisation of information manipulation, adapted from military doctrine, which this report uses to structure the full picture of situational awareness and FIMI understanding. Some of the levels describe an entity (political, strategic), while others focus on a (semi) permanent technical infrastructure (infrastructural, technical), or a time-bound, tailored campaign (operational, tactical). Figure 1 places the seven elements of the response framework at the level each belongs to.

[13] VIGINUM (SGDSN), Information Manipulation Set (IMS) — concept definition report, 22 January 2026.

D2.4 · FIGURE 1

Six levels, seven elements

Where each element sits, and what it costs the actor to replace



This project has received funding from the European Union's Horizon Europe Framework Programme under grant agreement N° 101132494.

SAUFEX

Figure 1. *The seven elements of the response framework placed at their level in the schematisation.*

2.9 Three IMSs, one Threat Actor

A Threat Actor can operate multiple Information Manipulation Sets, as the SDA case demonstrates. This report therefore adopts the decomposition set out in the January 2026 common operational picture of FIMI (a working group facilitated by EU DisinfoLab and endorsed by CheckFirst, VIGINUM, Cassini, the Auswärtiges Amt, the EEAS and DFRLab) which was applied independently by VIGINUM in its IMS concept report and by the EEAS in its fourth threat report.

Under that decomposition SDA's activity resolves into three distinct IMSs separated by tactic rather than by target:

1. **Doppelgänger,** which impersonates legitimate media through clone and typosquatted domains;
2. **RRN / Media Brands,** which builds inauthentic media brands of its own on custom platforms; and
3. **Undercut,** which is video-first and relies on AI voiceover.

The working group report is explicit that Media Brands/RRN "is considered a distinct IMS due to differences observed at the tactical level", and treats Undercut likewise. Retaining a single primary IMS would contradict the conceptualization of the deliverable, which aims at mapping out the full infrastructure of a threat actor to verify what countermeasures were effective to take it down and what failed. The set of response measures varies depending on the nature of the techniques observed: impersonating a real outlet may be unlawful, whereas merely establishing an inauthentic one is generally not. Therefore, the two tactics require different reporting routes and different legal bases to be addressed.

It should be noted that the activities and infrastructures reported from the May 2026 leak (physical false-flag operations, the SNG Media group, the fake think-tank and AI-ingestion projects planned for 2026) are examined as emerging sets, but are not counted as IMSs until the clandestinity and coordination criteria are independently and fully verified.

The IMS concept is interoperable with the CTI standards already central to this project: it is modelled on the CTI "Intrusion Set",[14] and designed for representation in STIX 2.1 and OpenCTI,[15] aligning this deliverable's methodology with the standardisation and

[14] VIGINUM (SGDSN), "Definitions and objectives of the concept of Information Manipulation Set (IMS)", 22 January 2026, https://www.sgdsn.gouv.fr/files/files/Publications/20260122_NP_TLP-CLEAR_SGDSN_VIGINUM_IMS_0.pdf.
[15] OpenCTI is an open source platform used to collect, store, organize, and visualize threat data using a knowledge graph based on STIX 2.1 standards.

automated-codification objectives of WP2. Its adoption by the EEAS, the G7 Rapid Response Mechanism and the organisations behind the common operational picture means the terminology recommended here is already in institutional use.

2.10 Cleaning the actor picture before counting

As described above, the raw codification was assembled with AI support under human supervision. Such a codification lists 1,085 distinct names across 1,838 mentions from 88 different sources, but that is the entire cast of items: threat actors, platforms, hosts, defenders, analyst sources and victims. Counting directly that raw set would inflate the figures and taint the attribution process. Cross-report resolution, which merges spelling and alias variants, reduces it to 880 distinct entities, and to 865 once the system of three IMSs is applied. Trust labels were assigned by name and applied uniformly across the whole corpus, so that no entity was considered harmful in one report and unknown or credible in the next.

When the technical and behavioural links required by the IMS coordination criterion were absent, adjacent manipulation sets were excluded from the SDA count, regardless of the cross-amplification of content between two IMSs. It is worth remembering that coordination relies on “strong technical or behavioural links among the assets, presumed operated by the same actor for the same sponsor”, while the other IMS criterion is clandestinity. For example, Pravda/Portal Kombat is a distinct IMS operated by a different actor, TigerWeb, which VIGINUM found shares no infrastructure with Doppelgänger despite documented mutual amplification.[16] CopyCop is likewise attributed to a separate operator, John Mark Dougan, working with the Centre for Geopolitical Expertise.[17] Whether CopyCop constitutes an IMS in its own right or a full component of Storm-1516 is still disputed in the literature, as cross-amplification can be a sign of mutual support between IMSs, but not a proof of membership. It is recorded in the registry as a relationship between sets, never as membership.

Cross-amplification should not be confused with coordination. When the bots of one set reshare links or content published by the bots of another, the link is visible to anyone and requires no common control, which is what separates it from the technical and behavioural links: shared hosting, a shared registrant, a common toolchain, or operational patterns that only a single operator could produce.

[16] VIGINUM, Portal Kombat technical reports, 12 February 2024, February 2024 (part two) and 28 April 2024. See also Centre for Countering Disinformation (Ukraine) with Molfar, Analytical Report: TigerWeb, August 2026,
[17] Recorded Future (Insikt Group), Russia-Linked CopyCop Uses LLMs to Weaponize Influence Content at Scale, 9 May 2024; U.S. Department of the Treasury, December 2024 (CGE/GRU designations).

Were cross-amplification treated as evidence of coordination, every set that reshares another would be absorbed into it, the chain would not end, and attribution would stop being meaningful. Figure 2 sets out the two links and which of them counts.

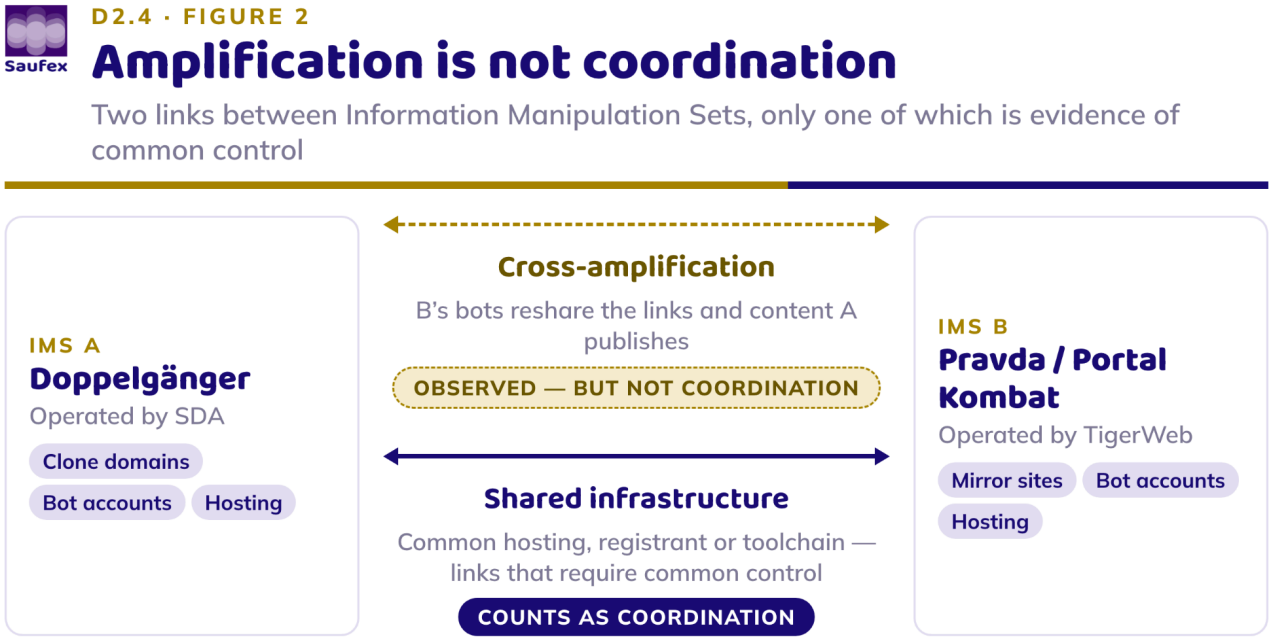


Figure 2. *Two kinds of link between Information Manipulation Sets.*

Cross-amplification is open to anyone and requires no common control, so it is not evidence that two sets belong together. VIGINUM found no shared infrastructure between Doppelgänger and Pravda / Portal Kombat despite documented mutual amplification, so the two are held as separate sets under separate actors.

The same discipline applies inside the SDA cluster, and it is what the evidence base supports: nineteen reports name Doppelgänger and RRN / Media Brands as separate entities, while ten (17% of the total) use a single name for both. The three SDA-operated IMSs are kept distinct from one another for the same reason they are kept distinct from adjacent IMSs: their techniques (fifth level, tactical) differ, and so does the response each one requires.

The report distinguishes the actor's operations that were only planned from those actually executed. Proposals and projects that emerged from internal documents are treated as plans. Any entry is counted at the operational level as part of an IMS infrastructure only when tactical-level activity (TTPs exploiting technical assets) was evidenced and reported.

After de-duplicating aliases and removing ecosystem noise, the proven cluster resolves to three core entities: the Social Design Agency, Structura National Technologies and ANO Dialog, together with a small number of front and subordinate structures. Structura is the company identified alongside SDA as responsible for the coordinated inauthentic behaviour directed at Germany, France, Italy, Ukraine and the United Kingdom, and designated with it by the European Union and the United States. ANO Dialog is the organisation outsourced to control institutional websites, to manipulate Yandex News results and to produce material on Ukraine. Both are set out more fully in Section 3. This is the Threat Actor. It operates three Information Manipulation Sets: Doppelgänger, RRN / Media Brands and Undercut.

The registry corroborates the cluster figures directly: SDA is named in 47 of the 88 reports, Structura in 32 and ANO Dialog in 21. The operations count is established in Section 5.2, where those evidenced as executed are separated from those recovered only as proposals. This cleaned picture is what the response record is then scored against.

The unit of analysis throughout is SDA's operation. A response measure is counted when it addresses an element of that operation, whether or not it names SDA: a designation naming several companies still reached SDA, and a provider's withdrawal of service reached the operation even where the source does not name the agency. Two measures in the registry were directed at ANO Dialog alone, the US designation of September 2024 and Microsoft's suspension of its email accounts. These are recorded but not counted. The footprint and the response are therefore measured against the same operation.

2.11 How the framework is scored against the record

The proposed framework defines seven elements on which the response can intervene, set out in Figure 3: clone domains; hosting and IP; social accounts and bots; ad-buying and payment; fake-brand content; AI ingestion; and the organisation itself. The level at which each element sits is shown in Figure 1. The recipient that can act on each, and the instrument available to it, are set out in Section 8.

D2.4 · FIGURE 3

The seven elements

What the response can intervene on, and what falls into each

1 Clone domains	Sites impersonating real outlets and institutions, including typosquats
2 Hosting and IP	The providers carrying the servers, and the prefixes they announce
3 Social accounts and bots	Accounts and automated networks that post and amplify
4 Ad-buying and payment	Advertising placements and the channels that pay for them
5 Fake-brand content	Content the operation produces, including material under invented outlet brands
6 AI ingestion	Material placed to be absorbed by models and search systems
7 The organisation itself	The companies, their officers and their contracting relationships

Figure 3. *The seven elements on which the response can intervene.*

To test the response against the SDA case, this report applies a two-column comparison to each of the seven elements:

- **Footprint**: what the operation controlled in that element (domains, accounts, hosting, ad and payment channels), counted only where the registry confirms it.
- **Response**: the binding measures delivered against that same element, counted as acts rather than as reports of them.

Since a single FIMI response or countermeasure can be described by many reports, and each description can differ in its wording, in the name given to the authority involved, and even in the date, if every retelling were counted, the response record would appear far larger than it is. Therefore, the methodological approach should be to count each act once, not each description of it, however differently the reports describe it. The September 2024 seizure of Doppelgänger domains shows how large the difference can be: this enforcement action appears in fourteen separate reports, which between them name the acting authority nine different ways and give five different dates.[18] If counted as separate items, it would inflate the response record roughly ninefold. That is why a consolidation procedure is key in this phase, as described in the above about cross-report entity resolution and event deduplication.

[18] Affidavit in Support of Seizure Warrant, United States v. Certain Domains, Case No. 24-mj-1395, United States District Court for the Eastern District of Pennsylvania, 4 September 2024.

There is a problem to measure what and how many response actions were actually taken vis-a-vis a certain FIMI asset and that concerns what should be considered in the comparison. When analysts publish a list of indicators of compromise alongside a report, they frequently continue updating that list for years afterwards. Such a list grows well beyond what the initial report itself contained at the publication date. For instance, Viginum's published RRN domain list contains 637 entries, of which 234 (37%) carry a creation date later than 13 June 2023, the day the report was first published. The most recent entry was registered in July 2025, more than two years after publication.[19]

Therefore, crediting the report with the whole list would attribute to it Indicators of Compromise (IOCs) that did not exist when it was written. For the purpose of this report, the footprint of a report only includes the IOCs that it contained when it was published, either because the indicator appears in the report, or because its first-seen date precedes publication. The full published list of IOCs, however, is retained separately, under its own provenance.

A domain enters the footprint only when a source attributes it to one of SDA's three Information Manipulation Sets. Several sources cover more than one operation, and the corpus includes whole domain lists for adjacent operations such as CopyCop and Portal Kombat. Counting every domain named in a source would have put 963 domains in the footprint, 445 of which belong to other operations or are not assets of the operation at all, such as the real outlets it impersonated. The footprint is therefore 518 domains, and the attribution of each is supplied with the registry.

It is possible that for several elements of the registry one column can be filled but not the other. For example, a report may establish what assets and techniques an operation was using while recording no response against it. Another report may record a response without establishing exactly the footprint that response addressed. In those cases the element is reported as unscored and no figure is estimated to fill the gap. The deliberate choice not to estimate missing numbers was made to guarantee each comparison rests on evidence that was codified from the reports.

In Part III a matrix of comparison with the scored elements is presented. To make an example of how that works, in the case of clone domains the registry attributes 518 domains to the operation, against which 8 hard actions such as seizures, takedowns and domain transfers were taken across the whole period, a ratio of roughly one action per 65 domains. It should be noted that the assets concentrated at the technical level are, by design, disposable. For instance, when in September 2024 authorities seized

[19] Authors' analysis of VIGINUM's published RRN indicator set, NDD_RRN.csv (637 domain names with registration dates), in VIGINUM (SGDSN), Rapports-Techniques, folder 202306 - RRN, https://github.com/VIGINUM-FR/Rapports-Techniques, accompanying the report first published on 13 June 2023 (English version "RRN: A complex and persistent information manipulation campaign", 19 July 2023).

Doppelganger domains, the threat actor and its infrastructure replaced them with new sites registered within 24 hours.[20] Instead, the durable levels were reached by other routes. The hosting that carries the operation was reached mainly by commercial providers withdrawing service and by the US and UK designations of Aeza, and the threat actors themselves by designations: the EU (2023), US (2024) and UK (2024) designations of the SDA cluster, and, for the sponsor level, the December 2024 EU listing of Zakharova.[21]

Whether the response reached the whole operation is therefore restated precisely: at which level of the IMS schematisation was each binding response delivered, and does the distribution of responses match the level at which the operation's continuity actually resides?

2.12 Two counts of response, deliberately distinguished

Part III reports two different counts of response, but they must not be read as competing, because they are deliberately distinct but complementary. The first one is about hard actions: an asset permanently removed, through a seizure, a deletion, a takedown or the transfer of a domain, or a person indicted. From the registry we can identify 8 of those hard actions against SDA, all against assets, since no indictment of SDA or its staff is recorded. A withdrawal of service is binding but not hard, since the asset moves to another provider. Whereas a binding measure is any legally or contractually compelling action, which adds sanctions, account suspensions, delistings, withdrawals of service and criminal investigations. There are 33 binding measures in the registry, of which 29 are dated, 3 are standing and 1 is undated.

The two counts answer different questions. The hard-action count asks how often the response went as far as it can go: an asset permanently taken, or a person charged. The binding-measure count asks how often enforcement of any kind was applied, whether or not anything was permanently removed, since a suspended account or a delisted domain can be restored and a sanctioned company continues to exist. Every hard action is also a binding measure, so the 8 are contained within the 33 rather than standing beside them.

2.13 Scope and limitations

The methodology of this report carries some acknowledged limits. Three follow from its design and one from the codification process. First of all, the reach of a FIMI operation and its actual outcome are not conflated, because the self-reported view

[20] E. Ponce de León, "Doppelganger websites persist one month following US government seizures", DFRLab, 9 October 2024.
[21] Council of the EU, Decision (CFSP) 2024/3174 and Implementing Regulation (EU) 2024/3188, 16 December 2024 (first listings under the hybrid-threats framework).

counts are treated as evidence that an internal KPI was claimed, not that a political effect was achieved. Oftentimes the only measurable data is the self-reported outreach, but not the real impact on the audience. In addition, the registry that was assembled from the codification process can only capture the residue of each operation: the artefacts that were detected and reported, but not what was missed from detection or removed by the threat actor before archiving. The same limit bears on the response. When a footprint is published, researchers release domain lists, account sets and indicator files. Those lists are the material from which this registry is built. However, a response is frequently not published at all. When a registrar cancels a domain, a hosting provider ends a contract or a payment processor closes an account, there is ordinarily no report, no announcement and no record an outside analyst could codify. The response column is therefore measured from a thinner source than the footprint column.

The third methodological limit concerns the scope of the analysis with a single case study. Since this framework was designed to test it against a real case, it has a value in showing where and why a response strategy fails, but cannot produce a statistical generalisation until it is replicated and tested with other cases. Furthermore, of the recorded 33 binding measures, 22 carry an effect assessment and 11 do not. The latter are not considered as evidence of a response failure, because it’s a gap of documentation rather than inaction.

I. The threat actor and its footprint

3. The threat actor in brief: SDA case analysis

3.1 SDA and the privatisation of Russian information operations

The Social Design Agency (SDA) is a Moscow-based commercial communications firm that conducts information manipulation for the Russian state under contract. It is not an intelligence service and not a state broadcaster. The United States Treasury describes it as a company "responsible for providing [the Government of Russia] with a variety of services" and acting "at the direction of the Russian Presidential Administration",[22] while the media investigation on the May 2026 leak describes it more plainly as "a Russian PR firm".[23] That commercial character is not a minor detail and matters for the whole argument. A state agency is always funded whereas an external contractor has to be paid and allocate resources for a specific project. A digital marketing agency performing tasks on behalf of the Russian apparatus has to buy the hosting and domains it uses, employ the right people and stay credible enough to win the next contract. Indeed, from the internal emails leaked in 2026 it appears that the previous leaks and failure have seriously undermined the trust of the Russian Presidential Administration into SDA's integrity and credibility.[24] Therefore, each of those dependencies is a point where response authorities can apply pressure: sanctions cut revenue, coercion on suppliers raises operating costs, naming employees individually makes the work personally risky, and public exposure damages the company's reputation. An intelligence service presents none of these vulnerabilities, while SDA offers the response framework more weaknesses to act on than a state body would.

The report examines the case study of SDA because it is at the same time the most demanding and well-documented case against which the framework can be tested. It

[22] U.S. Department of the Treasury (OFAC), designation of the Social Design Agency and Company Group Structura LLC, 20 March 2024.
[23] VSquare / Delfi / FRONTSTORY.PL, investigation into the May 2026 leak of SDA internal files, May 2026; OCCRP and Delfi, "Leaked Documents Reveal Russian 'Cognitive Strikes' Against the West", 24 May 2026.
[24] Gnida Project, ""Recent Social Design Agency leak points to Storm-1516 coordination with the Presidential Administration", 30 May 2026, https://gnidaproject.substack.com/p/sda-leak-pa-gru?utm_campaign=post-expanded-share&utm_medium=web.

is among the most thoroughly exposed FIMI actors on record. The formal attribution made by France, Germany, the United Kingdom and the United States, the designation for sanctions made by the European Union, the United States and the United Kingdom, a court-authorised seizure of its domains and two separate leaks of internal documents offer a rich pool of evidence. Despite the information collected, it is nevertheless still operating. The case is demanding for precisely this reason: when a threat actor and its infrastructure are poorly documented, a weak response can be explained by a lack of information, but that is not the case for SDA. Consequently, the explanation of the partial response failure must be different here. In Part III, the report examines the response record against the framework on that basis.

However, SDA appears to be part of a broader trend of privatisation and outsourcing, involving Moscow-based private firms that work for the Russian government on FIMI campaigns. Several media investigations and analyses have documented the links and partnerships between at least six PR agencies and IT companies that in various ways support the Russian propaganda strategy on the digital space, with clandestine operations or overt efforts relying on foreign influencers.

These companies work in parallel with the units of the intelligence services tasked with FIMI and active measures. Along with SDA, Structura National Technologies was already identified by media investigations and Meta, which deems it responsible for targeting Germany, France, Italy, Ukraine and the UK with a Coordinated Inauthentic Behavior (CIB).[25] A joint investigation by Meduza, Important Stories, and The Bell has uncovered the role of ANO Dialog, tasked with controlling institutional websites for internal propaganda purposes as well as with manipulating the Yandex News results and producing disinformation on Ukraine.[26] The first director of ANO Dialog was Alexey Goreslavsky, former deputy head of the Presidential Administration's Directorate for Public Projects, who later became the head of the Internet Development Institute, tasked with spreading propaganda in Russia and promoting the war in Ukraine.[27] ISD has identified a further Moscow-based company, ArgonLabs LLC, as working on Doppelgänger. It appears in four of the 88 sources and is recorded in the registry as related to the operation rather than as part of it, on the same basis as the adjacent sets discussed in the Methodology.[28] Lastly, the PR agency Limitless (Безграничные) has gathered a roster of international influencers (now removed from the website but

[25] Ben Nimmo and David Agranovich, "Removing Coordinated Inauthentic Behavior From China and Russia", META, 27 September 2022.
[26] Meduza, Important Stories, and The Bell, "Познакомьтесь с АНО «Диалог». Именно эта организация отвечает за пиар Минобороны РФ и создание фейков про Украину", 18 September 2023.
[27] Meduza, "«Чистыми руками взять патриотических денег и не зашквариться»", 6 June 2023.
[28] Elise Thomas, "Russian influence operation Doppelganger linked to fringe advertising company", ISD, 13 June 2024.

available in the archived version).[29] Among them is Russian propagandist Alexandra Jost aka “Sasha Meets Russia” paid by RT. The Limitless agency has received grants from the Presidential Foundation for Cultural Initiatives (PFKI), created by the Kremlin in 2021 to fund its propaganda efforts under the guise of cultural promotion.[30]
SDA remains the organisation we know more about, thanks to its prominent role and the reporting available, but the community of FIMI defenders should acknowledge the growing role of private and hybrid entities, because they offer new weak spots to be exploited by the response measures. Figure 4 sets out the state bodies and the private companies involved.

[29] Archived page, https://web.archive.org/web/20250919023412/https:/bezgranichnye.ru/bloggers, Wayback Machine.
[30] Jess Daly and Charles Smith, “How the Kremlin Buys Foreign Love Using Its Western Influencers Like Sasha Meets Russia”, United24media, 17 October 2025.

D2.4 · FIGURE 4

Who directs and who executes

Russian intelligence services, Kremlin coordination structures and the private companies contracted to run information operations

INTELLIGENCE SERVICES

GRU
Military intelligence
- Unit 55111 — Command of Information Operation Troops
- Unit 54777 — 72nd Special Centre
- Units 35555; 03126; 03127; 03128; 03132; 03134; 03138[9]
- African Initiative (Africa Corps) ●

FSB
Domestic intelligence
- 5th Service
- Unit 71330 — 16th Centre, also known as Berserk Bear

SVR
Foreign intelligence
- Unit 54939
- Africa Politology, also known as The Company[8] ●

↑ TASKING

KREMLIN

Special Influence Committees
Coordinated from the Presidential Administration by the Deputy Chief of Staff.[1]

Internet Research Agency
Dissolved in 2023.
- Wagner Group

● Their Africa operations passed to the units marked above.

Supported by the Presidential Foundation for Cultural Initiatives, created in 2021.

OUTSOURCING ↓

IT COMPANIES AND PR AGENCIES

Social Design Agency 2

Structura National Technologies 3

ANO Dialog 4

Argon Labs 5

Internet Development Institute 6

Limitless 7

SOURCES

1. PISM, "Russia Intensifies Disinformation Offensive Against Ukraine", 27 March 2024.
2. Radio Free Europe, "Leaked Records Detail Vast Russian Influence Campaign Targeting Ukraine, EU", 17 September 2024.
3. Meta, "Removing Coordinated Inauthentic Behavior From China and Russia", 27 September 2022.
4. Meduza, "Познакомьтесь с АНО «Диалог»", 18 September 2023.
5. ISD, "Russian influence operation Doppelganger linked to fringe advertising company", 13 June 2024.
6. Meduza, "Чистыми руками взять патриотических денег и не зашкваритьсяˮ, 6 June 2023.
7. United24, "How the Kremlin Buys Foreign Love Using Its Western Influencers", 17 October 2025.
8. Forbidden Stories, "Propaganda Machine", 14 February 2026.
9. CheckFirst, "Unveiling GRU's Information Operations Troops with OSINT and Medals", 9 February 2026.

Figure 4. *The Russian state bodies and private companies behind SDA's activity.*

3.2 The threat actor at a glance

The figures in this section are drawn from the cross-report registry described in the Methodology. Entities are resolved across all 88 sources, so that spelling and alias variants of the same organisation, person or asset are counted once. The corroboration column shows the number of corpus sources in which an item appears. That figure measures how widely an item is reported, not how independently it is confirmed. 13 of the 88 sources draw on leaked material, and several draw on the same leak, since the tranches obtained by Delfi Estonia were distributed to partner publications and one source in the corpus is an explicit republication of another. Where an item is known only from leaked material, a high count therefore reflects the breadth of reporting on a single body of documents rather than independent verification. Items resting on one or two sources are identified as weakly corroborated rather than presented in the same position as the rest. Table 3 covers the threat actor itself, not the wider ecosystem.

Table 3. The SDA cluster as the registry resolves it.

Component	As the registry resolves it, with the number of sources
Core contracting entities	Social Design Agency (47 of 88), Structura National Technologies (32), ANO Dialog (21)
Front and subordinate structures	ANO Dialog Regions, SNG Media, the internal cell "Center S" (one to two sources each)
Named individuals	Gambashidze (17), Tupikin (9), Tabak (5); eleven individuals resolve in total
The sponsor	Russian Presidential Administration: Kiriyenko (11), Zakharova (7)
Information Manipulation Sets	Doppelgänger (57), RRN / Media Brands (32), Undercut (6)
Evidenced domains	518 attributed to SDA's sets, of 963 named in a source body; a further 295 appear only in a dataset annex

The front and subordinate structures rest on one or two sources each and are weakly corroborated, as are the individuals below the three listed. Kiriyenko is identified in the

FBI affidavit as commissioning and supervising the activity, and Zakharova was designated by the European Union in December 2024. The counts for the Information Manipulation Sets are of sources naming each set. The narrower counts of sources attributing each set to SDA are given in Section 3.3.

3.3 Three Information Manipulation Sets, one Threat Actor

Following the taxonomy set out in the Methodology, SDA's activity is not treated as a single operation. It resolves into three Information Manipulation Sets (Doppelgänger, RRN / Media Brands, Undercut), distinguished by the techniques each relies on rather than by the country each targets. The distinction has practical consequences: impersonating an existing newspaper and founding an inauthentic media have different legal situations and are reported to different recipients. Table 4 gives the three sets and the sources attributing each to SDA.

Table 4. The three Information Manipulation Sets.

IMS	Sources attributing it to SDA	What distinguishes it
Doppelgänger	27	Impersonation of real media and government bodies through cloned and typosquatted domains
RRN / Media Brands	5	Inauthentic outlets under its own invented brands, presented as independent media
Undercut	4	Video-first production at volume with synthetic voiceover

Two features of these sets bear on what can be counted. Doppelgänger traffic is filtered through a commercial traffic-distribution system that serves investigators and targeted readers on different pages, which is why evidence has to be captured while an asset is live. RRN / Media Brands has been renamed twice, from Reliable Russian News to Reliable Recent News in 2022 and to Researchers and Reporters Network in December 2025, which accounts for several of the aliases the registry resolves for it. Undercut is operated alongside the other two sets rather than separately from them.

The choice follows the January 2026 common operational picture of FIMI, which treats RRN / Media Brands as “a distinct IMS due to differences observed at the tactical level”,

together with VIGINUM's IMS concept report and the fourth EEAS threat report.[31] In addition, nineteen sources from the corpus name Doppelgänger and RRN / Media Brands as separate entities.

What binds the three IMSs to a single threat actor is source attribution. Each is attributed to SDA by named sources, although the infrastructure indicators in the corpus are recorded predominantly under the Doppelgänger label rather than distributed across the three.

3.4 How SDA operates

By analyzing the corpus of sources, three recurring characteristics about how SDA operates can be identified. It works as an agency and the leaked material shows its activity is organised in client projects with budgets, monthly quotas and reach reporting. Even its internal measure of success is the ordinary KPI of a marketing agency: what audience was reached by a campaign and what was the share of voice and visibility of its content. Whether such reach turned into actual political effects is a separate question that the leaked documents do not address. That is why the methodology's two-axis provenance model keeps the actor's own reach claims flagged as claims rather than findings and evidence. Another important aspect of SDA's modus operandi is its reliance on commercial infrastructure with disposable assets. The IMSs depend on commercial traffic-distribution software, mainstream registrars and bulletproof hosting providers whose business model is to ignore abuse reports.

None of that infrastructure was built for disinformation, or used only for it. Kehr, the service that filtered visitors to Doppelgänger's clone sites, also served dating scams, and Aeza, its bulletproof host, serves ordinary cybercriminals as well.[32] The criminal use gives firmer legal ground than the political content does. Publishing false material is rarely unlawful, whereas hosting ransomware is, so a provider can be reached on grounds unconnected to information manipulation. The United States designated Aeza in July 2025 as a bulletproof hosting service enabling cybercriminals and technology theft, without mentioning disinformation, and the effect on SDA's hosting was the same.[33] A shared provider is also a shared chokepoint, since action against it reaches several unrelated actors at once.

[31] Doppelgänger Working Group, Building a common operational picture of FIMI, EU DisinfoLab, 15 January 2026; EEAS, 4th Report on Foreign Information Manipulation and Interference Threats, March 2026.
[32] Qurium, "When Kehr meets VexTrio", 13 November 2024, https://www.qurium.org/forensics/when-kehr-meets-vextrio.
[33] OFAC, Treasury Sanctions Global Bulletproof Hosting Service Enabling Cybercriminals and Technology Theft, 1 July 2025, https://home.treasury.gov/news/press-releases/sb0185.

The pressure point therefore lies in the supply of hosting and domains rather than in the artefacts they produce. Removing a clone site removes an output; removing the ability to buy hosting removes a capability.

Furthermore, the agency is increasingly combining online content with staged physical action (performed by proxies recruited by the Russian intelligence services). For example, in September 2025 some operatives placed severed pig heads with President Macron's name outside nine mosques and Islamic cultural centres in the Paris area, an act designed to spark religious hatred, xenophobia and islamophobia in France.[34] Three Serbian nationals pleaded guilty in Serbia in March 2026 to espionage, racial discrimination and criminal association and were sentenced to house arrest of between six and eighteen months.[35] The proceedings referred to training by an unnamed foreign intelligence service, therefore neither Russia nor SDA was named. Yet, it is the only criminal conviction in the corpus, and it reached the recruited proxies rather than anyone who commissioned the operation.

3.5 What the response has achieved, in brief

From the analysis of the reaction to SDA's activities, detection and attribution worked in most cases. The cluster has been named by France, Germany, the United Kingdom and the United States, designated in three jurisdictions, and documented across the 88 sources of the corpus.

The main issue was with what followed in terms of response, which was thinner. Binding measures were few and concentrated in a single year, with more than half in 2024, while all three Information Manipulation Sets continued to operate throughout. When the corpus shows what a measure achieved, degradation is the common result rather than a complete stop, which is rare.

Part III sets out the specific figures element by element. Section 12 follows one provider through the whole response to show where the limits of this approach lie.

[34] M. Laine, M. Vunš, and I. Lozovsky, "Leaked Russian Docs Reveal 'Cognitive Strikes' on the West", VSquare, 27 May 2026.
[35] ANews / AFP, "Serbian trio plead guilty over pig heads at French mosques", 6 March 2026, https://www.anews.com.tr/europe/2026/03/06/serbian-trio-plead-guilty-over-pig-heads-at-french-mosques.

4. How SDA came to be documented

4.1 The existing literature

The first public detection of SDA's activity dates back to September 2022, when the Doppelgänger campaign was simultaneously reported by Meta, EU DisinfoLab and Qurium Media Foundation.[36] In December 2022 Meta attributed it to two Russian private companies: Structura National Technologies and the Social Design Agency.[37] Attribution to the Doppelgänger cluster was therefore established within three months of first exposure.

Between 2023 and 2025, both technical and governmental documents increased public knowledge about this actor. In June 2023 the French agency VIGINUM published its technical report on the RRN campaign, six weeks before the first European Union designations.[38] The German Federal Foreign Office published its own analysis of German-language activity in June 2024.[39] In September 2024 the FBI released an affidavit supporting the seizure of 32 domains, which placed the SDA activity on the court record.[40] In parallel, researchers from NGOs and private companies mapped and documented the infrastructure used by SDA in its operations. Qurium shed light on the cloaking service, Recorded Future on obfuscation techniques and the IMS named Undercut.[41]

Finally, in 2026 there was a consolidation in the taxonomy related to SDA operations. VIGINUM introduced its definition of the Information Manipulation Set in January 2026.[42] The same month, a working group facilitated by EU DisinfoLab published its common operational picture of FIMI using IMS to strengthen technical attribution and disruption.[43] This was followed in March 2026 by the release of the fourth EEAS report

[36] Meta, "Removing Coordinated Inauthentic Behavior From China and Russia", 27 September 2022; EU DisinfoLab, "Doppelganger: Media clones serving Russian propaganda", 27 September 2022.
[37] Meta, "Recapping Our 2022 Coordinated Inauthentic Behavior Enforcements", 15 December 2022.
[38] VIGINUM, "RRN : une campagne numérique de manipulation de l'information complexe et persistante", 13 June 2023, updated 19 June 2023, English version "RRN: A complex and persistent information manipulation campaign", 19 July 2023; Council of the EU, "Information manipulation in Russia's war of aggression against Ukraine: EU lists seven individuals and five entities", 28 July 2023.
[39] Auswärtiges Amt, "Technischer Bericht: Desinformationskampagne Doppelgänger", June 2024.
[40] US DoJ, "Justice Department Disrupts Covert Russian Government-Sponsored Foreign Malign Influence Operation", 4 September 2024.
[41] Qurium, "When Kehr meets VexTrio", 13 November 2024; Recorded Future, "Obfuscation and AI Content in the Russian Influence Network "Doppelgänger" Signals Evolving Tactics", 5 December 2023; Recorded Future, "Operation Undercut Shows Multifaceted Nature of SDA's Influence Operations", 26 November 2024.
[42] VIGINUM, "Definitions and objectives of the concept of "Information Manipulation Set" (IMS)", Version 1.0, January 2026.
[43] EU DisinfoLab, "Building a common operational picture of FIMI", 15 January 2026.

on FIMI and in May the G7 Rapid Response Mechanism statement on a common understanding of the IMS framework, which both contributed to the consolidation of the taxonomy and conceptual schematisation.[44]

The existing literature has made substantial progress on attribution and taxonomy. What it has not produced is a response strategy: a statement of who should act on which part of an operation, with what instrument and to what evidential standard. That is the gap this deliverable sets out to address.

A second gap that is not addressed here concerns what an operation did to any audience. No source in the corpus establishes that a campaign changed what readers believed or their voting behavior. This deliverable establishes no such thing either. The only effectiveness this deliverable aims to address and measure is against the infrastructure of an operation.

4.2 Internal leaks and extracted data

Two leaks of SDA's own working documents changed the kind of evidence available.

Both were obtained by Delfi Estonia, the first together with Süddeutsche Zeitung, NDR and WDR, and reached the public through partner publications. The leaked files themselves are not sources in this deliverable, only the reporting on them is. Thirteen of the 88 sources draw on leaked material and the registry records what those publications established rather than the documents behind them. The reasons for that, together with the application of data protection rules, are set out in the Methodology.

The first leak, reported in September 2024 by VSquare, RFE/RL and other media partners, contained the organisational chart, the tasking chain to the Presidential Administration, several project files and budgets, and internal reports in which the agency graded its own campaigns by audience reached.[45]

The second leak, published in May 2026 by OCCRP and Delfi, documented activity that was not possible to attribute using open source material, such as the staged physical actions and AI-ingestion projects planned for 2026.[46] This leak also named SDA

[44] G7 Rapid Response Mechanism Position Statement on a Common Understanding of the Information Manipulation Set ("IMS") Framework, 15 May 2026.
[45] RFE/RL (Schemes), "Leaked Records Detail Vast Russian Influence Campaign Targeting Ukraine, EU", 17 September 2024
[46] OCCRP and VSquare, "Leaked Documents Reveal Russian 'Cognitive Strikes' Against the West — Including Islamophobic 'Pig Head' Attacks in Paris", 24 May 2026.

personnel, and the UK government designated 49 staff members for sanctions in the same month.

The analytical significance of the leaks derives from the insights regarding the organisational structure of SDA as well as from the confirmation about the intent of the threat actor, reinforcing the attribution of some actions. At the same time, they also introduce the risk the two-axis provenance model was theorised to manage: the reach figures claimed by the agency's reports to its client (the political sponsor) are treated as claims or goals rather than actual results of the operations.

5. Operations: process and record

This section deals with SDA's operations from two directions. The first describes the sequence any operation follows, from when it was commissioned to the account rendered to the client. The second section counts the operations themselves, separating those evidenced as carried out from those detected only as proposals.

5.1 How an SDA operation runs

This subsection describes the sequence of stages that an operation follows, from the point at which it is commissioned to the point at which it is reported to the client (the sponsor). The account is drawn from the 521 techniques codified across the 88 sources of the corpus, each assigned to the stage it predominantly serves. Figure 5 gives the six stages with the number of sources documenting each. Several techniques span more than one stage, so the counts show where documentation concentrates rather than a strict partition of the cycle.



Figure 5. *The six stages of an SDA operation, with the number of corpus sources documenting each.*

The assignment follows a written rule, supplied with the registry, which places a technique by the function it performs in the cycle rather than by the object it touches: a domain is acquired at the second stage, published through at the third and fourth and redirected through to conceal a destination at the fifth stage. The rule settled 349 of the 521 records with the assistance of AI and referred 172 to an analyst. The analyst reviewed 44 of these, comprising every record that could bear on the first or the last stage and every record the rule could not read at all. The remaining 128 are boundary cases among the middle four stages and resolving them would move records between stages that already stand within twelve of one another. Seven records are recorded as unstaged, three of them high-level narrative codes that describe what content is intended to do rather than how an operation runs.

The distribution of techniques among stages is itself a finding. The four middle stages appear in between 54 and 66 of the 88 sources, while tasking and planning and measurement and reporting appear in twelve each. The corpus therefore describes the middle of the cycle nearly five times more densely than either end. The reason is not that the ends matter less but that they are harder to see. The traces an operation leaves in acquiring assets, producing content, distributing it and evading removal are observable from outside, whereas the instruction that begins an operation and the assessment returned to the sponsor exist only in internal files. What is known about the first and last stages comes almost entirely from the reporting on the two leaks described in Section 4, not from technical analysis.

This matters for the response strategy because the central part of the operation cycle, which is the best documented, is also the part the threat actor can replace more easily. For example, stages 2 to 5 consist of assets that are bought and rented. Buying them again would be considered a cost rather than an operational obstacle. It would be way more expensive and costly to replace the elements at the beginning of the cycle (a political sponsor outsourcing the effort and tasking the agency for a FIMI operation). If the countermeasures such as personal sanctions, criminal investigations and law enforcement were more effective, it would be easier to dissuade the threat actor and its sponsor to initiate a new operation. Same goes for the ending part of the cycle: if the operation was successfully disrupted and mitigated, the threat actor cannot claim high performance results with its sponsor, who would be less inclined to task the agency again for further operations.

Two more aspects should be considered when analysing this cycle. The speed of the response determines whether it has any effect at all. The record does not measure the timing of an operation from commissioning to launch, since the tasking stage is barely documented, but does it with the time of the rebuilding. After the seizures of September 2024, replacement websites appeared on new top-level domains within

twenty-four hours and Meta records the operation moving spoofed sites to new hosting and new domains in less than a day.[47] The production side runs at the same tempo: cartoons were created and placed as advertising within forty-eight hours of the Hamas attack of October 2023 and of the Crocus City Hall attack in March 2024.[48] A countermeasure that takes weeks to deliver is therefore answering a configuration the operation has already abandoned.

If the cloned domains are taken down or the hosting providers are blocked only after a significant delay following detection, the response is useless against an operation that has already run and reached the end of its cycle. Also, thirteen sources document the traffic-distribution system, which uses obfuscation, filters and diverts visitors to different pages based on their geographical origin. This makes it more difficult to collect evidence and that is why this part of the operation should be tackled too by the response actions.

Finally, the stage of content production is evolving. Generative tools appear in 38 sources, principally for content, translation and synthetic voiceover, and their use is now routine rather than experimental. A smaller group of 8 sources documents a further step, in which the produced material is being ingested by LLMs and search systems rather than read by human users. While the impact of that technique is still to be determined, it should be noted that no countermeasure against it is recorded.

5.2 What was executed and what was only planned

Subsection 5.1 describes the cycle that an SDA operation usually follows. This one counts the operations themselves. Following the rule set out in the Methodology, an entry is counted as executed only where tactical activity was evidenced and reported. Proposals and projects recovered from internal files that were described by any source in the corpus are counted only as plans. Table 5 gives the result.

[47] E. Ponce de León, “Doppelganger websites persist one month following US government seizures”, DFRLab, 9 October 2024, https://dfrlab.org/2024/10/09/doppelganger-websites-persist; Meta, Adversarial Threat Report, Third Quarter 2024, December 2024.
[48] Check First, Reset.tech and AI Forensics, “Influence by Design: Doppelgänger, Sanctions, Meta and the Social Design Agency”, 17 January 2025, https://checkfirst.network/influence-by-design-how-meta-accepted-russian-propaganda-payments-despite-sanctions.

Table 5. Operations evidenced as executed and recovered as proposals.

Status	Count	Operations
Executed	13	Pig's Head; Green Synagogues; Armenian Genocide stickers; "Be greener!" car vandalism; Media-case Marseille; Zelenskyy Burj Khalifa hoax; Global South; Alley of Angels; CCI; CCE; CCU; The Other Ukraine; Good Old USA
Planned only	11	De Gaulle statue desecration; Seine sex-doll operation; Mitteleuropa / Vienna Agreement; erevan.one Armenia operation; Project 2026; Normal Israel; Baltic project office; Israel Outreach Project; International Conflict Incitement; Guerilla Media Campaign; US Influencers Network
Undetermined	4	Confederation of Independence; Project Magadan; Project RCA Media; Project Lipetsk

On that basis the corpus resolves to 28 distinct SDA operations: 13 executed, 11 planned but not evidenced as carried out, and 4 that cannot be assigned on the available evidence. Operations should not be confused with the Information Manipulation Sets, that are the standing apparatus for operations. The three identified sets (Doppelgänger, RRN / Media Brands, Undercut) are therefore not counted as operations.

The combination of different types of sources was instrumental in better understanding the nature and context of some operations. Because the detection of IOCs and technical observation can establish that an infrastructure exists, but it does not establish what an operation's rationale was or what it was launched for. This dimension of the case study, therefore, rests on the leaks, but the leaked files were not consulted directly. Each operation that was identified and counted has been recorded only through the sources from the corpus. The raw files were never taken into account, also in compliance with GDPR.

It is worth noting that the character of the executed operations has shifted over time. The earlier operations were conducted through the media and digital infrastructure only. The operations documented from the May 2026 leak are staged physical actions intended as false flags, including the severed pig heads outside Paris mosques, green paint at Holocaust memorial sites, and vandalism of German cars attributed to environmental activists. The earlier work sought audience share, whereas the latter sought a political incident for destabilisation. They also generated a criminal record, which the online activity generally does not.

In parallel, the set of planned operations matters for a different reason. It shows what the threat actor intended next and the strategic direction looks consistent: an alternative information ecosystem rather than more clone domains. Project 2026 lists a strategic studies think tank, an automatically maintained knowledge base of more than 200,000 pages aimed at German audiences, and a database of opinion leaders. The gap between the two sets is small in operational terms. The proposals require staff, budget and hosting infrastructures, all of which SDA already had, and none require capabilities it has not already demonstrated. A plan therefore becomes an operation at low marginal cost.

II. The response framework

6. Rationale for a multi-report approach

Part II sets out the framework this deliverable proposes for responding to FIMI operations.

We introduce the concept of multi-report approach, which is the practice of reporting each detected FIMI operation to all the stakeholders involved: the registrar of its domains, the provider hosting them, the platforms carrying its accounts and its advertising, the advertising and payment intermediaries through which it is funded, the regulators with jurisdiction over those intermediaries, and the authorities able to act against the organisation itself with sanctions or criminal proceedings. The multi-report approach entails reporting the operation again each time those elements are repopulated by the threat actor and its IMS.

An information manipulation operation is not a single object under a single jurisdiction. Section 3.3 established that SDA's three identified Information Manipulation Sets rest on different techniques: impersonating an existing newspaper and establishing an inauthentic one raise different legal questions, engage different remedies, and are therefore addressed to different recipients. That is why a multi-report approach is needed. No recipient has jurisdiction over more than a fraction of an operation, and a request addressed to one of them reaches only that fraction.

The registry records 33 binding measures against SDA, of which 2 are recorded as having stopped the activity they targeted and 5 as having prompted the operation to adapt rather than cease. Recorded action is also spread unevenly across the seven elements, and unevenly across those who took it. Hosting and IP drew ten binding measures, second only to clone domains, but all of them from commercial providers or from US and UK designations. Advertising and payment drew three, all from one platform, and AI ingestion none. These are the three elements Section 8 identifies as the hardest to reach.

Section 10 examines that distribution in full. What it shows is that the response so far has not been a multi-report response: it has acted on a few elements repeatedly and on others almost not at all, while all three Information Manipulation Sets continued to

operate. That establishes the problem rather than the remedy. The failure of a partial response does not by itself show that a concurrent one would succeed. It is possible that both fail, but the case for concurrency rests on two other grounds. No recipient controls more than a fraction of an operation, so a response addressed to one of them cannot reach the rest, regardless of how often it is repeated. Secondly, it is the single occasion in the record when several elements were addressed at once, between July and November 2024, whose recorded effects are the strongest in the corpus. In conclusion, what the record supports is that concurrency degraded the operation more than anything else attempted, not that it would have stopped it.

A further feature of the record argues for distributing reporting among recipients rather than simply increasing its volume. Almost everything the corpus records about platform enforcement comes from one company's own reporting and Section 7 explains why that cannot be read as a measure of what the platforms actually did. It is nonetheless a problem for a defender, because if the response really is concentrated on one company, it is exposed to any change in that company's policy, resourcing or commercial interest, none of which the community of FIMI defenders controls. If instead the other platforms acted but did not report it, a defender has no way to know where a report will be acted on. In any case, voluntary enforcement by platforms or service providers is an uncertain instrument. In Section 11 we show that the advertising intermediary continued to be paid by the operation after two designations.

As mentioned above, a FIMI operation often depends on commercial and disposable assets, which can be replaced as an ordinary expense rather than a long-term obstacle. A report that removes a domain therefore imposes only the cost of registering another, which is a routine stage of the operating cycle rather than an exceptional response to enforcement. An operation is a continuous process while a report is a single act in a specific moment, requiring repetition.

Hosting is not equally replaceable: the activity depends on the small number of providers prepared to disregard abuse reports and that constraint is what makes the hosting element the durable one. Therefore, a report that reaches the domain and the hosting at the same time imposes a cost of a different order from one that reaches the domain alone.

Reports must therefore reach the different recipients simultaneously rather than one after another. The analysis shows that sent in sequence, each removal can be repaired before the next arrives, and the operation continues uninterrupted.

Two legitimate objections can follow and the framework has to answer them. The first is that instruments differ by orders of magnitude in what they require. For example, an

abuse report to a registrar needs a domain and a screenshot, while a sanctions referral needs an evidentiary dossier and a criminal report takes months. Synchronising across that range would mean holding the fast instruments until the slowest was ready, by which time the assets they addressed have been replaced. Concurrency therefore applies within a class of instrument rather than across all of them. The routes that share a low evidentiary threshold and can be filed by any party are the ones to be brought together. Instead the slower instruments run at their own pace and are joined to them by a common evidence base rather than a common deadline, which is the argument of Section 9.

The second objection that can be made is that no party is in a position to coordinate the rest. The reports come from actors with no relationship to one another. For instance, in the sequence of 2024 two research organisations approached the providers directly and the four domain recoveries were filed by four brand owners acting independently. A researcher cannot tell a newspaper when to file, nor a treasury department when to designate. They may not even be able to tell that they are acting against the same Information Manipulation Set. The framework therefore does not assume a coordinator. What it requires is that each recipient can recognise the assets reported to it as belonging to one operation, which is what the IMS construct is for, and that reporting is cheap enough that convergence does not depend on anyone being persuaded. The proposed shared identifier does not yet exist and creating it is the first thing the framework asks for.

From these considerations it can be concluded that recipients must be able to recognise that the reported assets belong to one operation, which is what the IMS framework provides. The evidence supplied must meet a standard on which the recipient can act. When the intended outcome is on a judicial level, the evidence should be collected and preserved in a way that survives any legal challenge. This aspect will be addressed in the section about evidence standards. In addition, it is important that the reporting process should be repeatable at low cost even by organisations with limited dedicated capacity, which raises a question of standardisation and automation.

7. The ecosystem inventory: targets, intermediaries and responders

Section 6 argued that a FIMI operation must be reported to every party that controls part of it. Those parties have to be identified before they can be reported to. The goal of this section is to set out who they are, as the registry resolves them across the 88 sources, and to show where the evidence base is strong enough to name them and where it is not. Table 6 sets out the parties as the registry resolves them.

Table 6. The ecosystem of the SDA case.

Parties	Distinct entities	Named in 3 or more sources
Impersonated brands and other targets	114	16
Hosting providers	54	4
Registrars	14	1
Platforms	15	3
Researchers and civil-society defenders	62	17

The third column measures how well attested an entity is for each category. It should be noted that government responders are not counted at this stage because the registry does not separate state bodies acting against the operation from those acting for it. Instead, the measures governments delivered are analysed in Part III.

The 114 impersonated brands and targets are mostly news organisations whose identity was imitated. The best attested are Der Spiegel, named in 31 sources, and Le Parisien and T-Online in 8 each.

These organisations matter to the response framework for a practical reason: they own the trademarks that were infringed, which gives them a role and influence that other reporters do not have. A registrar or a hosting provider acts on a complaint from a brand owner more readily than on a complaint about political content. When the United States seized 32 domains in September 2024, the warrant was sought on money laundering and trademark offences, not on the ground that the content was

false.[49]

The table names 15 platforms that appear in the corpus, but one in particular dominates the record: Meta is named in 32 sources, against 5 for X and 4 for TikTok, so most of what is known about platform activity comes from Meta's own threat reporting. This poses an analytical challenge because when a platform appears less in the corpus it is unclear whether it was less involved in an operation or simply published less about it, and those two explanations point reporting effort in opposite directions.

On the other hand, researchers and civil-society defenders have a broader exposure with 62 entities, seventeen of them in three or more sources. The most frequently named are EU DisinfoLab, VIGINUM and Qurium. Just over half of the 88 sources come from this community rather than from governments or platforms, which makes it the largest single contributor to what is known about the threat actor of this case study.

The intermediaries used in the FIMI operations are the weakest part of the inventory. The registry identifies 54 hosting providers, four of which are named in three or more sources, and 14 registrars, of which only Namecheap reaches that threshold. The remaining thirteen rest on one or two reports each. Therefore, the reporting matrix in Section 8 cannot be organised around named providers, because for all but a handful the corpus cannot identify them reliably. The matrix has to be organised around the type of intermediary, with the provider verified case by case at the time of reporting.

[49] United States v. Certain Domains, Case No. 24-mj-1395 (E.D. Pa.), 4 September 2024

8. The reporting matrix

8.1 The matrix

Previous sections have established that a FIMI operation must be reported to several parties within the same period and those parties should be identified. This section sets out what each of them should receive. The matrix has two axes: the seven elements defined in the Methodology, and the type of recipient that can act on each. Section 7 showed why the second axis has to be a type rather than a specific company: the corpus names 14 registrars and 54 hosting providers but attests almost none of them well enough to build a framework around. The provider should therefore be identified case by case. Figure 6 gives the matrix.

D2.4 · FIGURE 6

The reporting matrix

Who must receive a report on each element, with what instrument, and what the record holds

Footprint
What the operation controlled in that element, counted only where the registry confirms it

Response
The binding measures delivered against that element, counted as acts rather than as reports of them

ELEMENT	RECIPIENT TYPE	INSTRUMENT	IN THE RECORD
1 **Clone domains**	Registrar, registry operator, brand owner	Abuse report; domain-dispute proceeding (UDRP or ccTLD)	**12 measures, 4 of them domain recoveries**
2 **Hosting and IP**	Hosting provider, upstream transit provider	Abuse report; sanctions referral	**10 measures, none by the EU**
3 **Social accounts and bots**	Platform	Notice and action; trusted flagger	**6 measures**
4 **Ad-buying and payment**	Advertising network, payment processor	Advertiser policy complaint	**3 measures, all by one platform**
5 **Fake-brand content**	Platform, search provider, AI and content-tool providers	Delisting; account termination	**1 measure**
6 **AI ingestion**	Model and search providers	Source exclusion	**None recorded**
7 **The organisation itself**	Regulator, law enforcement	Sanctions referral; criminal complaint	**5 designations**

Each instrument is set out in Sections 8.2 to 8.6. Counts are binding measures, as in Table 8.

Figure 6. *Recipient type and instrument for each element, with what the record holds.*

8.2 The bulletproof-hosting problem

As Section 3.4 stressed, part of the operation is hosted by providers whose commercial proposition is that they do not act on abuse reports. They are often referred to as bulletproof hosting, which are specialized web infrastructure services that deliberately ignore abuse complaints and takedown requests. Therefore, the ordinary instrument suggested in the matrix above (abuse report) has no effect on them. A report that a mainstream provider would act on produces no result with the bulletproof hosting providers, regardless of the evidence, because ignoring abuse reports is part of the service they sell and precisely what the threat actor pays them for. They offer the IMS operational resilience to countermeasures.

In this case, two courses of action are left and neither uses the abuse-report channel. The first would involve the upstream transit providers (IP transit providers) that are used by the uncooperative hosting company to reach the rest of the internet. These providers are usually mainstream companies with a public reputation to preserve and a business policy that includes abuse report procedures. These transit providers can be identified via public routing data. The second option is centered around the hosting providers, but focuses on designation with sanction referral that would hit them through the financial system, cutting them off from banks and payment processors in the jurisdiction, rather than asking them to take down a malign website through an abuse report. Both options have been used in this case study, but not in the order that might be expected. An upstream transit provider withdrew connectivity from Aeza in July 2024, eleven months before the first designation. Section 12 examines what each of the two measures achieved.

However, both courses of action suggested here require additional technical work before they can be used at all. The corpus names 54 hosting providers but attests only four of them in three or more sources, so in most cases the provider, and the upstream through which it connects, has to be verified at the time of reporting rather than taken from the existing sources.

8.3 Viable options and intervention toolkits

A comparison of the recorded response measures with the instruments available shows an inverse relationship. The instrument with the strongest legal basis, the brand-owner domain-dispute route, is the least recorded. The most recorded, leaving aside sanctions, which reach the organisation rather than its assets, is platform enforcement, which is also the most easily circumvented.

The brand-owner option appears four times in the registry. These are administrative proceedings rather than court actions. Under the Uniform Domain-Name

Dispute-Resolution Policy (UDRP), a trademark owner can obtain the transfer of a domain by showing that it is confusingly similar to a mark it holds, that the registrant has no legitimate interest in it, and that it was registered and used in bad faith.

Disputes are decided by accredited providers, such as the World Intellectual Property Organization (WIPO), a UN agency whose Arbitration and Mediation Center is the principal accredited provider administering UDRP cases. The outcome of the dispute binds the registrar. Country-code registries such as the Association Française pour le Nommage Internet en Coopération (AFNIC) run equivalent procedures of their own. Eligibility is not what limits the use of this instrument, since any owner of an infringed brand or domain can bring a claim. What limits it is that each proceeding concerns a single domain or a small set, while FIMI operations deploy thousands of them, and even when domains are recovered they are replaced quickly.

Platform enforcement is effective against precisely what it enumerates. The assets created to replace what it removes fall outside the measure, so each round requires a fresh one.

Two further options appear to be underused. Both Google News, the aggregation service, and Discover, the personalised feed, are shown in the Google app and on Android devices. They surface content to users who have not searched for it. This countermeasure can withdraw algorithmic recommendation but leaves the domain reachable through ordinary search, so it reduces distribution without removing the asset. Its advantage is that it can be applied at a scale that other measures, such as domain-dispute resolution, do not reach.

Connectivity and service withdrawal is the other option. Because it is the route the record shows achieving most, it is treated separately in Section 8.4.

AI ingestion was not addressed by any response measure, mainly because there is no specific instrument for it. A defender can tell a registrar that a domain infringes a mark or notify a platform that an account is inauthentic. In each case there is a defined recipient, a defined procedure and a defined remedy. For AI ingestion there is no equivalent for asking a model or search provider to exclude a known manipulation source: no established notice procedure, no criteria for what would qualify and no obligation to respond.

8.4 Service withdrawal on contractual grounds

Between July and November 2024 a sequence of commercial providers withdrew service from the operation, one after another: a hosting company, an upstream transit provider, a traffic-distribution service, a domain registrar and telecommunications carriers. The recorded effects of that sequence are the strongest in the corpus. It is the only occasion in the record on which several elements were addressed within a single period and the only one that forced the operation to rebuild across several elements rather than substitute one.

What distinguishes it is the basis on which each provider acted. None responded to a court order, a designation or a regulatory instruction. Each acted for breach of its own terms of service, which is a contractual matter between a provider and its customer and requires no finding that the content was unlawful or false. Aeza International closed two of its own bulletproof subsidiaries on 24 July 2024 and a week later Datapacket ceased to carry its traffic, after which the announced prefixes became unreachable within days. It is worth noting that the sequence of withdrawals was driven by researchers, not by governments. Correctiv and Qurium approached the providers directly, on the strength of their own technical work. No party in the sequence was obliged to act.

The route traced by this case should be reproducible thanks to the following factors. The recipient is identifiable from public data: which upstream network carries a given provider can be established by anyone from routing information that is free and continuously updated. The claim is framed as a breach of the provider's own terms rather than as a political judgement, which places it within the provider's ordinary commercial practice rather than asking it to adjudicate a dispute about content. Finally, the evidence required is what the provider needs to verify the breach, not what a court would need, so it can be assembled without the collection standard set out in Section 9, though evidence gathered to that standard serves both. This response route has limits too. It has no effect on providers whose commercial proposition is that they do not act on abuse reports, which is the problem set out in Section 8.2. In addition, its effects are reversible: Aeza redeployed behind new front proxies and by 2026 it was announcing more address blocks than before. In conclusion, it depends entirely on the willingness of a commercial party, which may change.

8.5 Standardisation of the reporting pipeline

The proposed matrix needs a common language and understanding descending from unified standards of reporting. The case study and more generally the history of FIMI response measures show that the reporting mechanisms are often disorganised and

episodic. These initiatives do not follow a strict protocol or a consolidated pipeline, also due to the different policies and procedures adopted by the organisations and companies involved. That is why standardisation matters and should be considered a priority to improve the quality and effectiveness of reporting. Reports are written independently of one another and several organisations may describe the same operation, each using its own format, naming the same assets and authorities in its own way, even dating the same events differently, to the extent that nothing in the reports themselves would show that they describe the same events.

The Methodology measures what the lack of standardisation can cost on the analytical side. For example, a single domain seizure is described fourteen times across the corpus, the acting authority is named nine different ways, and five different dates are given for the same action. A registrar or a platform receiving reports from several organisations faces the same difficulty in recognizing the assets as part of a single operation and often it can't treat them as one case.

That lack of standardisation can be addressed in three ways. The first is a shared identifier for each operation, recommended by the January 2026 working group on Doppelgänger, which is what the route in Section 8.4 depends on when more than one organisation reports the same operation. The second is a structured data format that the recipient's own system can ingest, such as STIX 2.1 as implemented in OpenCTI, with DISARM supplying the taxonomy for techniques. The third is a fixed report template for each recipient type, so that one operation can be reported to a registrar, a hosting provider and a regulator without being rewritten each time.

8.6 The trusted-flagger model

Article 22 of the Digital Services Act obliges online platforms to process notices from designated trusted flaggers with priority and without undue delay. The response bottleneck is not identifying assets but having reports acted on quickly enough to matter against an operation cycle that can run in days.

The mechanism is, in this case, a proposal rather than an observed practice. No trusted-flagger notice appears anywhere in the 320 recorded measures. The only Digital Services Act items in the registry are the formal proceedings the European Commission opened in 2024 against X, Facebook and Instagram, which are enforcement actions taken by a regulator against platforms for their own conduct rather than a channel through which a defender can report an operation.

The deliverable therefore cannot test the trusted-flagger mechanism against the SDA case and two features of the case explain why it does not appear in the record. One is

that most of what a defender would notify is not unlawful, which follows from the nature of FIMI itself: what can be reported is the coordinated inauthenticity of the network, in breach of the platform policy rather than EU law. The other concerns the form of the evidence submitted. A platform must give priority to a trusted-flagger notice, but the notice has to carry its evidence in a form the recipient can act on without investigating afresh. The measures that succeeded in this case did so on contractual grounds rather than by court order and the same logic could be applied to platforms whether or not the content is unlawful.

9. Evidence standards for court-admissible action

The recipients identified in Section 7 do not all require the same kind of proof to respond and implement countermeasures. A hosting provider acting on its own terms of service needs less evidence than a court does and a registrar deciding a domain dispute needs something different again. This proof asymmetry creates a practical challenge, because evidence gathered to satisfy the least demanding recipient cannot afterwards be escalated to the most demanding one. The working assumption of this framework is therefore that evidence should be collected once, to the highest standard any recipient might require, to be then presented in whatever form each recipient accepts. Table 7 sets out what each recipient needs established.

Table 7. Evidence standards by recipient.

Recipient	Basis for acting	What the evidence must establish
Platform, hosting provider, registrar	Breach of its own terms of service	That the asset is what it is said to be, and which term it breaches
Domain-dispute panel	Violations of trademark rights	Confusing similarity to a mark held by the complainant, absence of legitimate interest, and bad-faith registration and use
Regulator, on a sanctions referral	Statutory designation criteria	Attribution of the asset to a named entity, and that entity's role in the activity
Court	The applicable standard of proof	The provenance of every artefact, an unbroken record of its handling, and that it has not been altered

There are a number of features that sometimes make the highest standard hard to meet. Some of the assets deployed in a FIMI operation are ephemeral and quick substitutes can be found. Twelve replacement websites appeared within 24 hours of the mentioned seizures of 4 September 2024. Replacement registration is a routine stage of the FIMI operating cycle rather than a reaction to policy enforcement. That makes evidence not captured at the right time generally not recoverable.

Furthermore, platforms and digital services do not always preserve the evidence. Platform archiving is defined by its own classification rather than by investigative

need. Meta archives an advertisement only if it classifies it as political or issue-based. EU DisinfoLab identified at least twelve of the operation's advertising campaigns running on the platform and only a small share of their ads was classified that way.[50] The rest later disappeared, leaving no record behind. Enforcement and archiving are therefore in tension when the measure that removes an asset also removes the proof of what it was.

Regarding the evidence collection process, it should be noted that traffic-distribution systems serve different content to different visitors. Therefore, a capture made from the wrong location or IP may show the innocuous page rather than the operation's one. Establishing what was served, to whom and from where is part of the evidence burden. The affidavit supporting the September 2024 seizures shows what meeting legal standards looks like. Its 277 pages record FBI agents locating and reviewing six articles published on Doppelgänger websites. This review was done from inside the Eastern District of Pennsylvania and each part of that had a purpose: visiting the articles establishes what was actually published and doing so from that district establishes that the material reached readers there, which is what gave the court jurisdiction to issue the warrant. The collection was therefore deliberate rather than incidental, made while the material was still live and from a place chosen for legal reasons. Those are the conditions that allowed the evidence to support a seizure rather than only a published finding.

A second requirement follows from the replacement problem. Because assets are replaced rather than abandoned, evidence has to link successors to predecessors, or each round of enforcement begins again from nothing. The corpus contains a worked demonstration: the twelve post-seizure sites documented by DFRLab were tied to the seized domains through shared analytics and tag-manager identifiers, shared content-management usernames, redirect chains and content migrated from the originals. Identifiers of this kind survive the loss of the asset itself and allow a later report to say that a new domain belongs to an operation already established.

Two limits of the present evidence base should be recorded. The first concerns published indicator sets. They are not snapshots of a moment but lists that their authors continue to maintain. The Methodology showed what that means in practice: VIGINUM's RRN list now holds 637 domains, 234 of which were registered after the report it accompanies was published. A list of this kind cannot be used to show what an operation controlled at a given date unless it is first filtered to that date. If these lists were used unfiltered, it would put assets into evidence that did not exist when the

[50] Reset Tech, "So We Meet Again, Botiful: Latent Fake Ad Network Persists on Meta Two Years Since Doppelganger", 12 November 2024; EU DisinfoLab, "Doppelganger: Media clones serving Russian propaganda", 27 September 2022.

FIMI incident in question took place. The second limitation concerns the existing literature. Across the 88 sources, the corpus contains almost no discussion of how findings were preserved or of whether the material behind them would withstand examination by a registrar's legal department, a dispute panel or a court. The framework set out here therefore rests on what the DOJ affidavit demonstrates about deliberate collection and on what the two-axis provenance model in the Methodology requires.

III. Testing the framework against the SDA case

10. The response record and the scale mismatch

Part III measures the proposed framework against the selected case study. This section sets out what the response to SDA consisted of, element by element, and compares it with the footprint recorded in Part I. Section 11 draws the findings from it, and Section 12 examines the hardest element in detail.

The registry holds in total 320 records of response measures. After the consolidation described in Section 2.4 these resolve to 33 binding measures, of which 8 are hard actions. That gap is due to the same enforcement retold across the corpus, and to records that on review described measures against other actors or no binding measure at all. Counting the measures at record level rather than as acts would inflate the response roughly tenfold.

The mismatch of scale is the first result of the test. Against a footprint of 518 domains attributed to the operation, the registry records 8 hard actions: one action for roughly every 65 domains. The ratio is a lower bound on the response, since measures taken quietly by commercial providers seldom enter the published record. Against three Information Manipulation Sets operating continuously since 2022, it records 33 binding measures, 18 of them falling within a single year. Against 114 impersonated brands and other targets, it records four domain-dispute proceedings. The response is also distributed unequally across the seven elements, as Table 8 shows.

Table 8. The response by element.

Element	Binding measures	Source reports
The organisation itself	5	46
Clone domains	12	33
Social accounts and bots	6	11

Element	Binding measures	Source reports
Hosting and IP	10	11
Ad-buying and payment	3	12
Fake-brand content	1	1
AI ingestion	0	0

The first column counts acts, the second the source reports that describe them. Three acts touch more than one element, since a single takedown can remove an account and the domain it pointed to, and they are counted under each, so the first column sums to 37 rather than 33. The assignment of acts to elements is supplied with the registry.

A single record can touch more than one element, since the same takedown can remove both an account and the domain it pointed to. The numbers therefore show where action is concentrated. A caution applies to the second column: it counts how often a measure was reported, not how much was done. As already stated, almost everything the corpus holds on platform enforcement comes from one company's own reporting, for the reason given in Section 7, so the column reflects who publishes as much as who acts.

Yet, two patterns are clear enough to survive. The first is a difference between what was done and what was reported. The five designations of the organisation are described by 46 source reports, about nine each, while the ten measures against hosting are described by 11, about one each. The part of the response that reached the middle of the operation is therefore the least documented, and an analyst reading the corpus by volume would see designations everywhere and hosting almost nowhere. The second pattern is about who acted. Every one of the ten recorded measures against hosting came from a commercial provider or from a US or UK designation. The one EU listing of a provider the operation had used, Stark Industries in May 2025, is not recorded in the corpus and is discussed in Section 14.2. Advertising and payment drew three measures, all from one platform, and AI ingestion none. These are the three elements Section 8 identified as the hardest to reach. The first because the provider is often paid to disregard abuse reports, the second because the recipient of the complaint is being paid by the operation and is in a conflict of interest, the third because no pre-defined reporting course of action exists at all.

The second mismatch concerns the type of entity the response targets. Almost every measure in the registry is addressed to a specific thing: a web domain, a social media account, a hosting arrangement, an advertising channel, a company or a person. Of the 320 records, only one names an SDA operation as its target and in that case it was about journalistic exposure rather than a binding measure. On the contrary, the threat actor does not work in those units but by operation, with internal project names and budgets. Section 3.4 records that its own measure of success is the reach that a campaign has achieved or claims to have achieved.

In conclusion, the response counts what assets it removed, while the threat actor counts what audiences it reached. The practical consequence is that the questions about countering operations cannot be answered from the record, because no measure was ever directed at either. Effectiveness can be reported only regarding the single assets that have been removed, which is the element that a FIMI operation can replace most cheaply.

The domain-dispute route produced the clearest outcomes. AFNIC returned leparisien[.]re to Le Parisien in February 2024. Süddeutsche Zeitung recovered sueddeutsche[.]me in November 2022 and Le Monde recovered lemonde[.]ltd in July 2023, both through UDRP proceedings. In October 2023 a WIPO panel ordered the registrar NameCheap to return gouv[.]fm to the French foreign ministry, whose identity the operation had impersonated at diplomatie[.]gouv[.]fm to share false statements. All four proceedings recovered the domain.[51] Four cases are too few to establish a success rate, but the figure that matters more is how rarely this instrument is used. It accounts for about 1% of the recorded response, against 114 brands and other targets recorded as impersonated. Impersonated outlets often chose a different channel: Ukrainian media went to the state rather than to the registrar. Both Obozrevatel and Ukrainska Pravda asked the Security Service of Ukraine to act, and RBC-Ukraine notified the National Police, the SBU and the Ministry of Digital Transformation. None of those referrals has a recorded outcome, while all four domain disputes recovered the domain.

The pace of countermeasures matters too. When the US Department of Justice seized 32 Doppelgänger domains on 4 September 2024, twelve replacement sites appeared on other top-level domains, most of them within twenty-four hours, carrying content migrated from the seized domains and tied to them by shared analytics identifiers. An instrument that works reliably against one asset cannot by itself answer an operation that replaces assets faster than proceedings can be brought.

[51] EU DisinfoLab, “Assessing cost-effectiveness: responses to the Doppelganger operation”, 24 May 2024.

Platform enforcement is the opposite case, being the most used instrument in the record. Account suspensions and takedowns account for 73 records, more than any element other than the organisation itself.

Domain blocking by Meta and X made links to the blocked domains unshareable, while blacklisting pushed the operation into the multi-stage redirect chains described in Section 5.1.[52]

Exclusion of domains from search and discovery worked at a different order of volume, being the one instrument in the record applied to hundreds of domains at a time rather than to individual assets. In February 2025 Google's Threat Analysis Group blocked 47 Doppelgänger domains from appearing on Google News and Discover. Then it blocked a further 235 domains during the same year belonging to Portal Kombat, an adjacent set excluded from the SDA count under the coordination criterion. Of those, only 47 belong to SDA. The larger figure is recorded here because it shows that the instrument scales in a way the others do not, having been applied to 235 domains of a single adjacent set in one year.

Between July and November 2024 a series of commercial providers withdrew service from SDA's infrastructure, one after another. It is the only sustained sequence of action against several elements at once that the record contains. Aeza International shut down two of its own bulletproof hosting subsidiaries, Lethost and Sunhost, on 24 July 2024, forcing the front-proxy infrastructure to move elsewhere. A week later Datapacket stopped providing upstream connectivity to Aeza and its associates, after which the announced autonomous system numbers disappeared.[53] In October the operator of the Kehr cloaking service blocked the operation and shut the service down, while a Ukrainian provider blocked the Russian customer following an inquiry. The recorded outcome of both was that most of the operation's links on X stopped working.

[52] Correctiv, "Inside Doppelganger – How Russia uses EU companies for its propaganda", 22 July 2024, https://correctiv.org/en/latest-stories/2024/07/22/inside-doppelganger-how-russia-uses-eu-companies-for-its-propaganda/; European External Action Service, "Doppelganger strikes back: FIMI activities in the context of the European elections 2024", June 2024, https://euvsdisinfo.eu/uploads/2024/06/EEAS-TechnicalReport-DoppelgangerEE24_June2024.pdf.
[53] Qurium Media Foundation, "When Kehr meets VexTrio", 13 November 2024, https://www.qurium.org/forensics/when-kehr-meets-vextrio/.

On 7 November 2024 the German registrar 1API GmbH blocked 28 domains "after a brief review due to violations of our terms of service", leaving a large proportion of the portals inaccessible,[54] and during the same month European telecommunications providers withdrew service from SDA and blocked Real Reliable News.[55]

Two features of this sequence matter for the response framework. It reached five different kinds of providers within four months: a hosting company, an upstream transit provider, a traffic-distribution service, a domain registrar and telecommunications carriers. Its recorded effects are the strongest in the corpus, since most of the operation's links on one platform ceased to function and a large share of its portals became inaccessible.

The operation then redeployed behind new front proxies and replaced its assets, so the sequence degraded the operation rather than halting it, but it remains the only occasion on which several elements were addressed at once, which is what Section 6 argues the response should do as a matter of course.

The advertising and payment element holds about a dozen records, nine of them concerning Meta's own advertising products. One frequently cited figure cannot be read as enforcement: of 98 pages used to publish 98 advertisements, only two remained online, but the source attributes the removals to either Meta or the operators themselves, so they cannot be counted as a response.[56] Of 509 advertiser pages identified in 2024, 372 were still online when Reset Tech reported in November.[57]

Check First, Reset.tech and AI Forensics estimate that Meta earned approximately $338,000 from advertising linked to SDA between August 2023 and October 2024, across more than 8,000 detected advertisements, and that payments continued after both the European Union designations of July 2023 and the United States designation

[54] Correctiv, "Doppelganger: CORRECTIV investigations bring Russian propaganda campaign to a halt", 15 November 2024, https://correctiv.org/en/fact-checking-en/2024/11/15/doppelganger-correctiv-investigations-bring-russian-propaganda-campaign-to-a-halt/; The Insider, "German investigative journalists confirm Doppelgaenger bot network's digital trails lead to Russian Ministry of Defense", 18 November 2024.
[55] J. Pamment and D. Tsurtsumia, Beyond Operation Doppelgänger: A Capability Assessment, Lund University and the Swedish Psychological Defence Research Institute, 12 May 2025, https://www.psychologicaldefence.lu.se/sites/psychologicaldefence.lu.se/files/2025-05/Beyond%20Operation%20Doppelg%C3%A4nger.pdf.
[56] Counter Disinformation Network / Alliance4Europe, Fool Me Once: Doppelganger Continues to Persist on X and Facebook, 1 September 2024, https://alliance4europe.eu/report/fool-me-once-russian-influence-operation-doppelganger-continues-on-x-and-facebook.
[57] Reset Tech, "So We Meet Again, Botiful: Latent Fake Ad Network Persists on Meta Two Years Since Doppelganger", 12 November 2024, https://www.reset.tech/resources/reset-tech-investigation-doppelganger.

of March 2024.[58] The platform responsible for enforcement was therefore being paid by the operation it was asked to act against. The European Commission opened formal proceedings against Meta under the Digital Services Act on 30 April 2024, on grounds including deceptive advertising.

[58] Check First, Reset.tech and AI Forensics, "Influence by Design: Doppelgänger, Sanctions, Meta and the Social Design Agency", 17 January 2025, https://checkfirst.network/influence-by-design-how-meta-accepted-russian-propaganda-payments-despite-sanctions.

11. Findings: successes and limitations

The first finding from the test of the framework against the SDA case study is that detection and attribution worked. The cluster has been named by France, Germany, the United Kingdom and the United States, designated by three jurisdictions, placed on a court record through the seizure affidavit, and documented across 88 published sources by governments, platforms and independent researchers. Its three Information Manipulation Sets have been separated and characterised, its infrastructure mapped to the level of individual providers and its internal organisation exposed by two leaks.

The second finding concerns what the measures achieved once taken. Of the 33 binding measures, the corpus records an assessed outcome for 22 but none for the remaining 11 (Figure 7). Among the 22, two are recorded as having stopped the activity they targeted, eleven as having degraded it without stopping it, five as having prompted the operation to adapt in order to continue, and four as having produced no change or as not possible to assess from what the source reported. The most common documented result is therefore degradation and complete cessation is the least common, in two cases. Eight of the 22 outcomes were coded by the analysts, for the reason given in Section 2.4, and five are contested between their sources. The 11 unassessed measures are reported here as unassessed rather than as failures, since the absence of a recorded outcome is a gap in the reporting and not evidence about the measure.

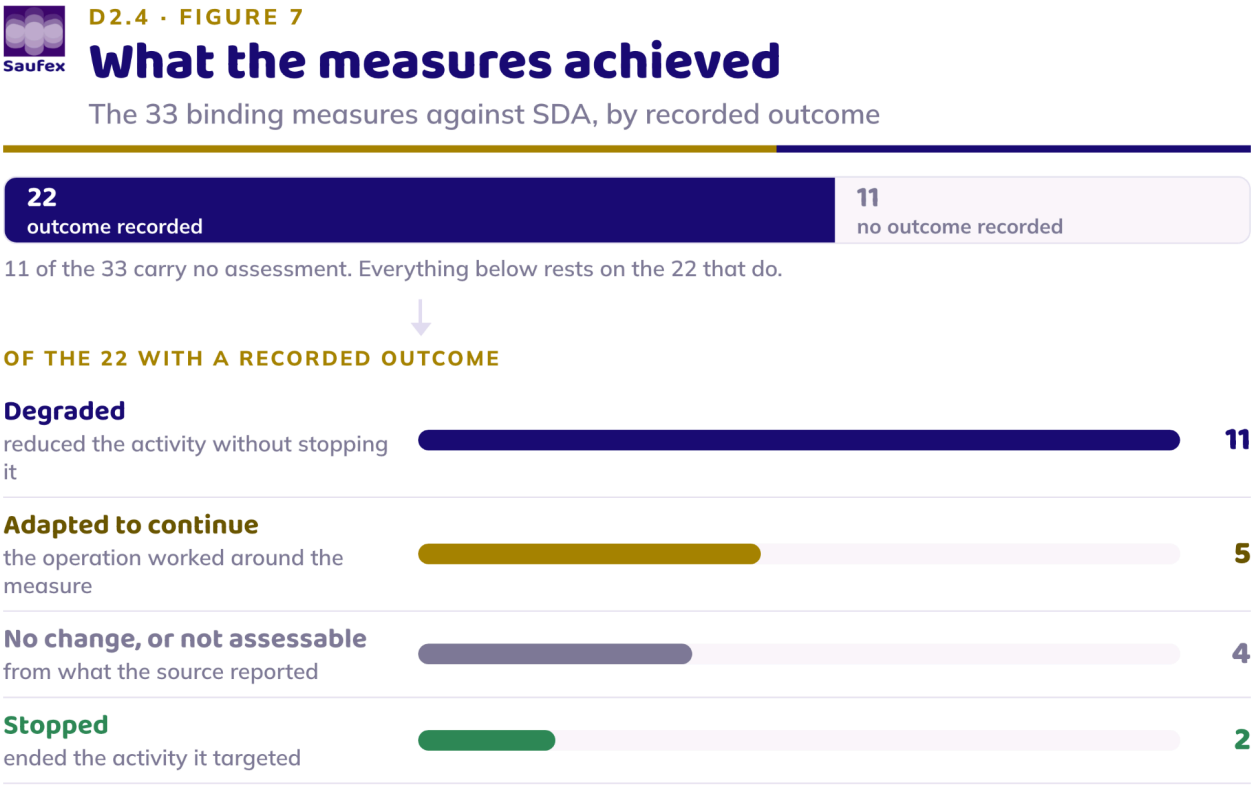


Figure 7. *The 33 binding measures against SDA by recorded outcome.*

The third finding is about which routes produced the recorded successes. The four proceedings that recovered a domain were brought by brand owners through administrative dispute procedures. The sequence with the strongest recorded effects, between July and November 2024, consisted of commercial providers withdrawing service on their own terms of service, at the prompting of two research organisations. Neither route involved a court, a designation or a regulator. Meanwhile the two instruments carrying the most institutional weight performed less well on the record: the September 2024 seizure of 32 domains was answered by twelve replacement sites within twenty-four hours, and the July 2025 designation of a hosting provider left it operating openly fourteen months later. This is not an argument that state instruments do not matter, since designation reaches levels no commercial route can touch. It is rather showing that the instruments most available to non-state defenders are also, on this record, among the most effective, and that they are the least recorded.

The fourth finding is that the principal failures are structural rather than procedural. The registry records 8 hard actions (seizures, takedowns and domain transfers) against 518 domains that the corpus attributes to the operation. That is one action for roughly every 65 domains. The recorded response understates what was done, for the reason given in the Methodology, but no plausible allowance for unreported measures closes

a gap of that order. Another issue concerns the mismatch of units. Every measure was addressed to an asset or to a company, none to an operation, so the record cannot say whether any of SDA's twenty-eight operations was stopped. In addition, the three elements Section 8 identified as hardest to reach were barely reached by EU measures: hosting through commercial withdrawals and US and UK designations, with a single EU listing of a provider the operation had used, advertising and payment through a single platform, and AI ingestion not at all. None of these is corrected by performing the same measures more diligently, because each requires a different addressee, a different legal basis, or an instrument that does not yet exist.

The fifth finding concerns advertising, where the problem differs in kind from the other elements. Unlike hosting where the provider is paid to disregard reports, in advertising the recipient was both reachable and responsive, but was also selling advertising to the operation and continued to be paid after two designations. A response framework that relies on the voluntary enforcement of an intermediary has to account for the possibility that the intermediary's commercial interest points the other way.

Two main limitations of this assessment should be recorded alongside the findings. It should be noted that a third of the binding measures carry no assessed outcome and that eight of the remaining outcomes were coded by the analysts, so every statement above about effect rests on 22 acts, that is a small base from which to generalise, and it is why this section reports proportions among the assessed measures rather than rates across the whole. Also, nothing in the corpus measures effect on audiences. The response record can say how many domains were seized and how many accounts suspended. But it cannot say whether any of it reduced the influence of the material on anyone and SDA's own reach figures are claims made to the client-sponsor rather than findings. Effectiveness throughout this deliverable therefore means effectiveness against infrastructure, not in the mitigation of the campaign towards the audience.

12. Countering the hardest element: sanctioning bulletproof hosting (the AEZA case)

This section focuses on a specific case study that helps illustrate the challenge to counter the exploitation of hosting providers, which is the element that resists reporting by design, and one that EU measures have so far barely reached. It is possible to follow one provider, Aeza, through the whole response to observe how the strongest available instrument was applied to the hardest available target.

On 24 July 2024, Aeza International shut down two of its own bulletproof hosting subsidiaries, Lethost and Sunhost, which forced the front-proxy infrastructure to move. A week later Datapacket stopped providing upstream connectivity to Aeza and its associates. A hosting company reaches the internet only if a larger network agrees to carry its traffic and announce its address ranges to everyone else. When Datapacket stopped doing that, the announcements ceased and the address ranges became unreachable. Neither act required a court, a regulator or an official designation. Both were commercial decisions taken after researchers approached the companies concerned.

In July 2025, instead, the response came from the state institutions. The US Treasury designated Aeza Group, three affiliated companies (Aeza International, Aeza Logistic and Cloud Solutions) and four of its senior figures, describing it as a bulletproof hosting service enabling cybercriminals. The United Kingdom designated Aeza International, the group's British company, in September 2025, and Aeza Group itself in November 2025 for providing bulletproof hosting to SDA.[59] On 19 November 2025 the United States, Australia and the United Kingdom acted jointly against the hosting provider Media Land, and the United States designated further companies and individuals that it identified as helping Aeza evade the sanctions, among them Hypercore and the Serbian company Smart Digital Ideas.

In the two weeks to 13 September 2026, fourteen months after the first designation, Aeza Group's autonomous system AS210644 was announcing 317 IPv4 prefixes and 41 IPv6 prefixes, and was reaching the rest of the internet through thirty upstream networks.[60]

[59] UK Sanctions List, Russia regime: Aeza International Ltd (RUS3052), designated 19 September 2025; Aeza Group LLC, designated 19 November 2025. The September 2025 designation is not recorded in any source in the corpus and is therefore not among the binding measures counted in Section 10.
[60] RIPE NCC, RIPEstat, "Announced Prefixes" and "AS Neighbours" data calls for resource AS210644 (AEZA-AS, Aeza Group LLC), retrieved 13 September 2026, covering the measurement window 30 August to 13 September 2026, https://stat.ripe.net/AS210644.

Three weeks earlier, on 25 August 2026, the same measurement returned 289 IPv4 prefixes.

After designations in two countries and a further round against its successor companies, the provider is still operating, and the second measurement returned more announced address space than the first, not less. The designations did not reach the provider because a designation works through the financial system: it bars persons within the designating jurisdiction from dealing with the target, and it is effective in proportion to how much of the target's business runs through that jurisdiction. Aeza's connectivity does not run through the jurisdictions that designated it. Its traffic reaches the wider internet principally through the service provider Aurologic GmbH, a German network operator, which appears in the routing data as its main transit provider by a wide margin.[61] Providing transit is an ordinary commercial service, but the point is jurisdictional. Aeza has been designated by the United States and by the United Kingdom, not by the European Union, so the route by which it reaches European audiences lies in the one jurisdiction that has taken no measure against it.

A single commercial decision by one transit provider achieved in a week a disruptive effect that the designations had not produced fourteen months later. The only measure that visibly interrupted the provider was indeed Datapacket's withdrawal of upstream connectivity in July 2024, after which the announced prefixes disappeared within days. This should not be considered the decisive instrument, because its effect was temporary and the operation redeployed behind new front proxies. However, while designation may be slow and bounded by jurisdiction, connectivity withdrawal reaches the service itself, immediately, and is bounded only by the willingness of the upstream to act.

In conclusion, designation and connectivity are complementary instruments rather than alternatives. In this case they were used eleven months apart by different actors without coordination, which reduced the effect of both.

[61] On RIPEstat's measure of how prominently each transit relationship is observed, aurologic GmbH scores 654 against 146 for the next provider and 87 for the one after that.

IV. Lessons learned and recommendations

13. Lessons learned

Testing the framework against the record yields six lessons. Two tools organise them. The first is the Pyramid of Pain, borrowed from cyber threat intelligence, which ranks the parts of an operation by how much it costs the actor to replace them. The second is the response loop: respond, monitor what the actor does next, and respond again to its adaptation.

13.1 Knowledge was not the constraint

The SDA cluster has been attributed by four governments, designated by three jurisdictions, placed on a court record and documented across 88 published sources. Its three Information Manipulation Sets have been separated, its infrastructure has been mapped down to individual providers, and two leaks have exposed its internal organisation. Nothing in the record suggests that the response fell short for want of knowing who was responsible or what they controlled. More investment in detection would answer a question this case has already settled.

13.2 The EU response barely reached the middle of the pyramid

David Bianco's Pyramid of Pain ranks the indicators defenders use against an adversary by how much it costs the adversary to lose them: hashes and IP addresses at the bottom, which are cheap to change, and tools and techniques at the top, which are expensive to change.[62] This report adapts the model from detection to response. Each level is ranked by how long the operation took to rebuild after an asset was removed, as the corpus records it, and each is matched with the parties able to act on it (Table 9).

[62] Bianco, D. J., The Pyramid of Pain, Enterprise Detection & Response, 1 March 2013, http://detect-respond.blogspot.com/2013/03/the-pyramid-of-pain.html.

Table 9. The Pyramid of Pain adapted to FIMI and applied to the SDA record.

Level	SDA example	Rebuilt after removal	Binding measures	Who can act
Sponsor	Russian Presidential Administration (Kiriyenko, Zakharova)	Not replaceable by the actor	Included with the organisation (Zakharova: EU, Dec 2024)	Governments: attribution, personal sanctions
Organisation	SDA, Structura, ANO Dialog	Still operating after designations in three jurisdictions	5	Sanctions authorities, prosecutors
Services and servers	Aeza hosting, Kehr cloaking, advertising and payment channels, pathways into AI systems	Slower; Aeza redeployed behind new front proxies and was still announcing 317 IPv4 prefixes in Sept 2026	13: hosting 10, advertising and payment 3, AI 0	Hosting and transit providers, CERTs, sanctions authorities, advertising and AI providers
Linking artefacts	Shared analytics and tag-manager IDs, CMS usernames, redirect chains	Survive the loss of a domain	Evidence only	Researchers and investigators, who tie new assets to old ones
Domains and accounts	518 domains, including clones of real outlets; bot and advertiser accounts	Within 24 hours (12 sites after the Sept 2024 seizure)	18: domains 12, accounts 6	Registrars, platforms, brand owners through domain disputes, courts
Content	Cloned articles, videos, advertisements	Hours (cartoons within 48 hours of the Oct 2023 attack)	1	Platforms, AI and content-tool providers

Source: *Adapted from Bianco, 2013; SAUFEX consortium, 2026.*

Read against the pyramid, the middle of the operation was reached, but hardly by the EU. Measures against the organisation account for 5 binding measures, and measures against domains and accounts for 18. Hosting, advertising and payment drew 13, and the pathway into AI systems none. Every one of the 13 came from a commercial provider, from a platform, or from a US or UK designation. The Union's only listing of a provider the operation had used, Stark Industries in May 2025, falls outside the corpus and is not counted. Yet the middle layer produced the strongest recorded effects: the withdrawals of hosting, transit and cloaking services, together with a registrar and

telecommunications carriers, between July and November 2024.

The lower levels still matter. A domain seizure or an account takedown costs the operation money and time, produces evidence, and pushes it into more expensive workarounds such as multi-stage redirect chains. What the pyramid shows is the price of that effort. At the bottom, a measure has to be repeated as often as the asset is replaced, which in this record meant within 24 hours. Higher up, each measure lasts longer, as long as the infrastructure below it is reached too.

13.3 Each measure was a single strike

An operation that adapts is contained only if someone watches what it does next and acts on it. The record shows the first step many times and the next two rarely. Of the 33 binding measures, 11 have no recorded outcome, so nobody reported what followed them. Of the 22 with an outcome, 5 led the operation to adapt and continue. After the seizure of 32 domains in September 2024, twelve replacement sites appeared, most of them within 24 hours, and researchers documented them in detail. The record does not show who, if anyone, acted on them.

The measures were also aimed at a unit that makes follow-up hard. Of 320 recorded response measures, one names an SDA operation as its target, and it is an act of journalistic exposure. The actor plans, budgets and reports by operation. Because no binding measure was aimed at an operation, the record cannot say whether any of SDA's 28 operations was stopped.

The one period in which pressure was sustained across several elements, from July to November 2024, produced the strongest effects in the corpus. It lasted four months, and each withdrawal followed the one before. What distinguished it was that the pressure continued after the first step.

13.4 What worked was decentralised

The measures with the clearest results were taken by parties outside the state. Brand owners recovered all four domains they disputed through the UDRP and AFNIC procedures. Hosting, transit, cloaking, registrar and telecommunications providers withdrew service under their own terms of service after Correctiv and Qurium approached them. Neither route required a court, a regulator or a designation. This is direct empirical support for the second SAUFEX objective, the decentralisation of response.

The same evidence shows what decentralisation lacks when it is left to itself. Four domain disputes against 114 impersonated brands and other targets means the route was almost unused. The 2024 sequence depended on two organisations working with project funding, with no shared identifier for the operation, no link to the designations that followed, and no record of outcomes.

13.5 State instruments were not carried through

Designation reaches what no commercial route can reach: the company, its officers and their access to finance. In this record it was not carried through to the infrastructure and services that the operation depends on. The advertising platform kept receiving payments from the operation after the EU designation of July 2023 and the US designation of March 2024, about USD 338,000 between August 2023 and October 2024 on the estimate of Check First, Reset.tech and AI Forensics.

The legal basis for stopping these sales existed twice over. From July 2023 SDA's assets were frozen in the EU and since 17 December 2022 Council Regulation (EU) No 833/2014 has prohibited providing, directly or indirectly, advertising, market research and public opinion polling services to any legal person established in Russia, whether or not it is listed (Article 5n(2a), inserted by Regulation (EU) 2022/2474, now Article 5n(1)(d) as amended by Regulation (EU) 2025/2033).[63] The prohibition binds every company incorporated in a Member State, including those through which the largest platforms sell advertising in Europe. It applies where the provider knew or had reasonable cause to suspect whom it was serving, which is where notification by researchers and authorities matters. The Commission's guidance on Article 5n does not address online advertising and no enforcement case under it has been published.

Aeza shows the same gap at the hosting level. The United States designated Aeza Group in July 2025 and the United Kingdom in November 2025. In September 2026 its network was still announcing 317 IPv4 prefixes through thirty upstream networks, and it reaches the internet mainly through Aurologic GmbH, a German transit provider. The route by which it reaches European users therefore runs through the one jurisdiction that has not designated it. The one measure that visibly interrupted the provider was commercial: Datapacket's withdrawal of transit in July 2024, eleven months before the first designation. The two instruments were never used together.

[63] Council Regulation (EU) No 833/2014 of 31 July 2014, Article 5n, consolidated text, https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:02014R0833-20260724; Council Regulation (EU) 2022/2474 of 16 December 2022; Council Regulation (EU) 2025/2033 of 23 October 2025; European Commission, "Frequently asked questions: provision of services", updated 22 January 2026, https://finance.ec.europa.eu/publications/provision-services_en.

13.6 A contractor can be hurt where a business is hurt

SDA is a commercial contractor. It has to buy hosting and domains, employ staff, pay suppliers and keep its client satisfied, and each of these is a pressure point that an intelligence service would not offer. The clearest damage at the top of the pyramid in the whole record came from exposure: reporting on the 2026 leak indicates that the Presidential Administration came to doubt SDA's reliability after the earlier leaks (Section 3). Institutions do not control leaks or journalism, but the lesson carries over. Measures that raise the actor's costs and damage its standing with its client reach the top of the pyramid, which takedowns alone do not.

14. Recommendations

14.1 The approach

The recommendations are addressed to institutions, as Task 2.9 requires, and the record defines their role. The routes that worked were decentralised, and institutions should make them systematic: fund them, give them a shared picture of the operation, and connect them to the instruments that only institutions hold. Each recommendation names the body that would act, the action and the instrument. They follow the response loop of Section 13: respond at every level of the pyramid, monitor what comes back, and counter-respond, supported by a common evidence standard and by safeguards. Table 10 summarises them, and Sections 14.2 to 14.10 give the evidence for each.

Table 10. Recommendations by addressee, action and instrument.

Recommendation	Who acts	What they should do	Instrument	Priority
1. Designate the infrastructure, in step with commercial action	Council of the EU; European Commission; national sanctions authorities	Extend designations from the actor to the services that carry it, starting with Aeza Group; notify transit and hosting providers when a designation is adopted	EU restrictive-measures regimes already applied to the SDA cluster	**High**
2. Stop payments from designated actors	National sanctions authorities; European Commission (DSA); Resilience Councils	Require advertising platforms and payment processors to screen advertisers and payers against sanctions lists and for establishment in Russia, and report matches	EU asset-freeze obligations; Regulation 833/2014, Article 5n(1)(d); DSA Articles 34, 35 and 39	**High**
3. Notify brand owners and providers	Resilience Councils; Digital Services Coordinators	Tell impersonated outlets that their domains can be recovered; send terms-of-service notices to hosting, transit, cloaking and registrar providers	UDRP and ccTLD dispute procedures; provider terms of service; report templates (Section 8.5)	**High**

Recommendation	Who acts	What they should do	Instrument	Priority
4. Fund and mandate monitoring after each measure	European Commission; Member States; Resilience Councils	Fund standing monitoring of what the operation rebuilds; require the outcome of every measure to be recorded	Multi-annual funding line; outcome reporting	**High**
5. Share measures and outcomes under one operation identifier	EEAS (Rapid Alert System); FIMI-ISAC; G7 Rapid Response Mechanism	Adopt a shared IMS identifier; record the operation each measure targets and what followed	STIX 2.1 course-of-actio n objects in OpenCTI; DISARM	**Medium**
6. Counter-respond at national level	National CERTs; Digital Services Coordinators	CERTs notify the upstream transit of designated or repeatedly reported hosts; DSCs work with trusted flaggers on network-level notices	Public routing data; DSA Article 22	**Medium**
7. Adopt a common evidence standard	Digital Services Coordinators; European Board for Digital Services	Adopt the Section 9 standard as the code of practice on evidence gathering; fund the preservation of captures	DSC guidance; Board guidelines; SAUFEX training material	**Medium**
8. Create a notice route against LLM grooming	EU AI Office; European Commission (DSA)	Define who receives a request to exclude a known manipulation source, what qualifies, and a duty to respond	AI Act obligations for general-purpos e AI models; DSA Articles 34 and 35 for very large online search engines	**Medium**

Source: *SAUFEX consortium, based on the registry evidence in Parts I to III.*

14.2 Designate the infrastructure, in step with commercial action

Designation is the one instrument that only states hold. In this record it reached the organisation and not the services that carry it. The European Union has designated the threat actor, its subsidiaries and eleven named individuals, but not the provider that carries the operation today. It has listed hosting providers before: Stark Industries under the hybrid-threats regime in May 2025, a provider to which the operation had moved part of its infrastructure in August 2024, and Media Land under the cyber

regime in July 2026.[64] Aeza Group, designated by the United States for enabling cybercrime and by the United Kingdom for hosting SDA campaigns, still operates, and its main transit provider is German. The Council and the Commission should consider extending the existing listings to the hosting and service providers that carry designated operations, starting with Aeza, as they have already done for two other hosting providers. When a designation is adopted, national sanctions authorities should notify the transit and hosting providers in their jurisdiction, so that the commercial withdrawal that acted fastest in 2024 follows a designation within days, not a year later.

14.3 Stop payments from designated actors

The advertising platform kept receiving payments from the operation after the EU and US designations: about USD 338,000 between August 2023 and October 2024, across more than 8,000 advertisements, on the estimate of Check First, Reset.tech and AI Forensics. Whether an advertiser is on a sanctions list can be checked without judging any content. National competent authorities for EU sanctions should require advertising platforms and payment processors to screen advertisers and payers against the consolidated sanctions lists and to report matches. Screening should also cover establishment in Russia, since Article 5n(1)(d) of Regulation (EU) No 833/2014 prohibits selling advertising to any company established there, listed or not. The Commission should address online advertising in its guidance on Article 5n. The Commission can address the same risk for very large online platforms through their risk-assessment and mitigation duties under Articles 34 and 35 of the Digital Services Act and the advertising repository under Article 39. The Description of Action gives Resilience Councils the task of undermining FIMI actors' remuneration and payment models, and they can flag suspected cases.

14.4 Notify brand owners and providers

The two routes with the clearest results in the record are both open to parties outside the state. Brand owners recovered every domain they disputed, yet only four disputes were filed against 114 impersonated brands and other targets. Providers withdrew hosting, transit, cloaking and registrar services under their own terms of service when

[64] Council of the EU, "Russian hybrid threats: EU lists further 21 individuals and 6 entities and introduces sectoral measures in response to destabilising activities against the EU, its member states and international partners", 20 May 2025, and Council Implementing Regulation (EU) 2025/965; Council of the EU, "Russian cyber-attacks and destabilising activities: Council sanctions nine individuals and four entities", 13 July 2026. On the move to Stark Industries: Qurium, "When Kehr meets VexTrio", 13 November 2024. Neither listing is recorded in the corpus, so neither is among the binding measures counted in Section 10.

researchers presented the evidence. Neither route needs a new law. What both lack is someone who tells the brand owner that a clone exists and can be recovered, and someone who sends the provider a notice it can act on. Resilience Councils, as the intermediaries between state and non-state actors described in D1.1, and Digital Services Coordinators are placed to do both, using the per-recipient report templates proposed in Section 8.5.

14.5 Fund and mandate monitoring after each measure

Repeating a report each time an asset returns requires someone to watch for its return. Today a few research organisations do that work, case by case and on project funding, and it is the step the record shows to be missing: 11 of the 33 binding measures have no recorded outcome, and seized domains were replaced within 24 hours. The European Commission and Member States should fund standing monitoring of what an operation rebuilds after each measure, covering replacement domains, moved hosting, returning accounts and continuing payments. Resilience Councils can coordinate that monitoring nationally. Every institution that takes a measure should record its outcome, so that the next responder knows what came back.

14.6 Share measures and outcomes under one operation identifier

Responders act at different speeds and in different jurisdictions, and they should keep doing so: a single country acting alone moves faster than a coalition, and the record shows national and commercial measures working without coordination. What they lack is a shared picture of which operation a measure targeted and what followed. Only one of the 320 recorded measures names the operation it targets. The EEAS, through the Rapid Alert System for Member States, FIMI-ISAC for civil society, and the G7 Rapid Response Mechanism should adopt a shared identifier for each Information Manipulation Set, building on the IMS definition already adopted by the EEAS and the G7 Rapid Response Mechanism, and record against it every measure and its outcome. Expressing measures as STIX 2.1 course-of-action objects in OpenCTI, with DISARM countermeasure codes, would let this picture travel in the format WP2 has standardised for analysis.

14.7 Counter-respond at national level

When an operation adapts, the next measure is often national. Finding which upstream network carries a hosting provider requires only public routing data, and the one measure that visibly interrupted Aeza was a transit provider's withdrawal. National CERTs should identify the upstream transit of hosting providers that are designated or repeatedly reported, and notify it. Digital Services Coordinators should work with the trusted flaggers they certify under Article 22 of the Digital Services Act, so that notices on impersonation and other unlawful conduct by coordinated networks cover the network as a whole and carry evidence the platform can act on without investigating afresh. No trusted-flagger notice appears among the 320 recorded measures.

14.8 Adopt a common evidence standard

Evidence collected to a platform's standard cannot later support a court, and an asset not captured while live is usually lost. The FBI affidavit behind the September 2024 seizures shows the standard that works: deliberate capture of live material, from a recorded location, with the handling of every artefact documented. Section 9 turns this into requirements, including the identifiers that tie replacement assets to the ones removed. The Description of Action foresees a code of practice for Digital Services Coordinators on evidence gathering, labelling, notice and takedown. Digital Services Coordinators and the European Board for Digital Services should adopt the Section 9 standard as that code and fund the preservation of captures that platforms do not archive. The SAUFEX training material for Digital Services Coordinators can carry the standard into practice.

14.9 Create a notice route against LLM grooming

SDA's content reaches people through its sites and videos. It also reaches AI systems and search engines, which absorb it and repeat it to users who never visit the source, a practice known as LLM grooming. Eight sources in the corpus document the operation's work in this direction, and its planning for 2026 includes a knowledge base of more than 200,000 pages aimed at German audiences, counted here as a plan. The registry records no measure against it. Removal does not help, because material already absorbed by a model or an index is reproduced after the source is gone. A defender can report a clone domain to a registrar, but no office exists to ask a model or search provider to exclude a known manipulation source, no criteria define what qualifies, and no provider has a duty to respond. The EU AI Office, through the AI Act obligations on general-purpose AI models with systemic risk, and the Commission,

through the risk-mitigation duties of very large online search engines under the Digital Services Act, can create these three components. Researchers can draft the criteria for the regulator to adopt. The planned projects give a reason to act before they run.

14.10 Safeguards

Several of these recommendations act without a court: withdrawal of service under terms of service, delisting and exclusion from AI systems. The Description of Action requires FIMI to be addressed without resorting to censorship and with respect for freedom of expression. Four conditions keep the measures proportionate.

- **Conduct, never viewpoint.** Action targets impersonation, coordinated inauthenticity or sanctions evasion, and never an opinion.
- **A narrow definition.** A manipulation source is an asset attributed to an identified Information Manipulation Set to the evidence standard of Section 9.
- **Transparency.** Notices and their grounds are recorded, and published once the measure is taken.
- **Redress.** The affected party can contest the measure, through the internal complaint and out-of-court dispute settlement routes the Digital Services Act provides for platforms, and through equivalent routes for other intermediaries.

15. Conclusion

This deliverable proposed a response framework and tested it against the public record of one threat actor, the Social Design Agency. The response to SDA did not fail for lack of knowledge. Four governments attributed the operation, three jurisdictions designated it, a US court authorised the seizure of its domains and 88 published sources document it. The actor still operates.

Read against the Pyramid of Pain, the EU response reached the top and the bottom of the operation and barely reached the middle. Designations reached the organisation. Takedowns reached domains and accounts, which the operation replaced within a day. Hosting, advertising and payment drew 13 binding measures, all from commercial providers, a platform, or US and UK designations, and the pathway into AI systems none. The platform that sold the operation its advertising kept being paid after two designations, on the estimate of Check First, Reset.tech and AI Forensics. The prohibition on selling advertising to companies established in Russia has been part of EU law since December 2022, but no guidance or published enforcement has yet applied it to online advertising. Aeza, its bulletproof host, was still announcing 317 IPv4 prefixes in September 2026, fourteen months after its US designation, and reaches the internet mainly through a German transit provider, in the one jurisdiction that has not designated it.

The response was also a series of single strikes. A third of the binding measures have no recorded outcome, and in 5 of the 22 cases with an outcome the operation adapted and continued. The record rarely shows anyone monitoring what the operation rebuilt, and no binding measure was aimed at an operation as such. The one period of sustained pressure, from July to November 2024, produced the strongest effects in the record.

The measures that worked were decentralised: brand owners recovering domains, and providers withdrawing service at the prompting of researchers. This is the evidence behind the second objective of SAUFEX. It also shows what decentralisation lacks when left to itself: funding, a shared picture of the operation, and a link to the instruments that only states hold.

The recommendations follow from these findings. Institutions should respond at every level of the pyramid, fund and mandate the monitoring of what comes back, and act on it, while each country and each responder keeps acting at its own speed. Each of the eight recommendations in Part IV names the body that would act: the Council and

the Commission on designations and payments, Resilience Councils and Digital Services Coordinators on notice and follow-up, national CERTs on transit, the EEAS and its partners on a shared operation identifier, and the AI Office on LLM grooming.

The evidence measures effects on infrastructure. No source measures effects on audiences, and every statement about outcomes rests on the 22 measures with a recorded result. A single case can show where a response fails, and it cannot show how often. The next step is to apply the same test to a second actor. Storm-1516, already present in this corpus as an adjacent set, is the obvious candidate.

Figures and Tables

List of Tables

Table 1. The controlled vocabulary of this deliverable
Table 2. Types of source in the corpus and what each establishes
Table 3. The SDA cluster as the registry resolves it
Table 4. The three Information Manipulation Sets
Table 5. Operations evidenced as executed and recovered as proposals
Table 6. The ecosystem of the SDA case
Table 7. Evidence standards by recipient
Table 8. The response by element
Table 9. The Pyramid of Pain adapted to FIMI and applied to the SDA record
Table 10. Recommendations by addressee, action and instrument

List of Figures

Figure 1. The seven elements of the response framework placed at their level in the schematisation
Figure 2. Two kinds of link between Information Manipulation Sets
Figure 3. The seven elements on which the response can intervene
Figure 4. The Russian state bodies and private companies behind SDA's activity
Figure 5. The six stages of an SDA operation, with the number of corpus sources documenting each
Figure 6. Recipient type and instrument for each element, with what the record holds
Figure 7. The 33 binding measures against SDA by recorded outcome

Annexes

Annex A - The list of the 88 sources

Sources are listed alphabetically by publisher. The identifier is the one used in the consolidated registry and in the per-report codifications, so that any figure in this deliverable can be traced to the sources behind it. Identifiers are registry keys and do not encode publication dates.

No.	Publisher	Title	Date	Type	Identifier
1	***Bellingcat (with Investico, Trouw, De Groene Amsterdammer)***	Using Open Source Internet Routing Tools to Monitor a Sanctioned Russian Bank	2024-03-01	Journalism	Bellingcat_routing
2	***Bloomberg News***	Leaked Files Show Russia's Plan to Influence AI and Search Results	2026-06-23	Journalism	Bloomberg_Project2026_2026-06
3	***Check First / AI Forensics / Reset Tech***	Influence by Design: Doppelgänger, Sanctions, Meta and the Social Design Agency	2025-01	Leak-derived	CheckFirst_InfluenceByDesign
4	***ClearSky Cyber Security***	Doppelgänger NG | Russian Cyberwarfare campaign	2024-02-22	Research / OSINT	ClearSky_DoppelgangerNG
5	***CORRECTIV***	Inside Doppelganger – How Russia uses EU companies for its propaganda	2024-07-22	Journalism	Correctiv_EUcompanies_2024-07
6	***CORRECTIV***	Doppelganger: CORRECTIV investigations bring Russian propaganda campaign to a halt	2024-11-15	Journalism	Correctiv_KehrDisruption_2024

No.	Publisher	Title	Date	Type	Identifier
7	***Council of the European Union***	Information manipulation in Russia's war of aggression against Ukraine: EU lists seven individuals and five entities	2023-07-28	Government	EUCouncil_2023sanctions
8	***Council of the European Union***	Council Decision (CFSP) 2024/2643 concerning restrictive measures in view of Russia's destabilising activities	2024-10-08	Government	EUCouncil_hybrid_framework
9	***Council of the European Union***	Russian hybrid threats: EU agrees first listings in response to destabilising activities against the EU, its member states and partners	2024-12-16	Government	EUCouncil_hybrid_listings
10	***Counter Disinformation Network (CDN) / Alliance4Europe***	Fool Me Once: Doppelganger Continues to Persist on X and Facebook	2024-09-01	Research / OSINT	Alliance4Europe_FoolMeOnce
11	***DFRLab (Atlantic Council Digital Forensic Research Lab)***	Explainer: the Russian influence operations targeting the 2024 US elections	2024-09-06	Research / OSINT	DFRLab_explainer
12	***DFRLab (Atlantic Council Digital Forensic Research Lab)***	Doppelganger: How Russia mimicked real news sites and created fake ones to target US audiences	2024-09-18	Research / OSINT	DFRLab_USelection
13	***DFRLab (Atlantic Council Digital Forensic Research Lab)***	Doppelganger websites persist one month following US government seizures	2024-10-09	Research / OSINT	DFRLab_persist

No.	Publisher	Title	Date	Type	Identifier
14	***EU DisinfoLab***	Doppelganger: Media clones serving Russian propaganda	2022-09-27	Research / OSINT	EUDisinfoLab_Doppelganger_2022
15	***EU DisinfoLab***	Assessing cost-effectiveness: responses to the Doppelganger operation	2024-05-24	Research / OSINT	EUDisinfoLab_CostEffectiveness_2024-05
16	***EU DisinfoLab (Doppelganger Working Group)***	Building a common operational picture of FIMI: Using IMS to strengthen technical attribution and disruption	2026-01-15	Research / OSINT	DGWG_common_operational_picture
17	***European External Action Service (EEAS)***	2nd EEAS Report on Foreign Information Manipulation and Interference Threats — A Framework for Networked Defence	2024-01-01	Government	EEAS_2nd_FIMI
18	***European External Action Service (EEAS) STRATCOM***	Doppelganger strikes back: FIMI activities in the context of the EE24	2024-06-01	Government	EEAS_DoppelgangerStrikesBack_EE24
19	***European External Action Service (EEAS), Policy Planning and Strategic Communication***	4th EEAS Report on Foreign Information Manipulation and Interference Threats: Dismantling the FIMI House of Cards	2026-03-01	Government	EEAS_4th_FIMI
20	***European External Action Service (EEAS), Strategic Communication and Foresight (SG.STRAT)***	3rd EEAS Report on Foreign Information Manipulation and Interference Threats - Exposing the architecture of FIMI operations	2025-03-01	Government	EEAS_3rd_FIMI

No.	Publisher	Title	Date	Type	Identifier
21	*FBI / Cyber National Mission Force (CNMF) with Netherlands AIVD, MIVD, DNP and Canada CCCS (Joint Cybersecurity Advisory AA24-191A)*	State-Sponsored Russian Media Leverages Meliorator Software for Foreign Malign Influence Activity	2024-07-09	Government	GT-0120_Meliorator
22	*Forbidden Stories*	Propaganda Machine: Secret documents reveal Russia's foreign influence strategy across three continents	2026-02-14	Journalism	ForbiddenStories_2026
23	*Foreign, Commonwealth & Development Office (GOV.UK)*	UK action against Russian foreign information warfare	2025-12-09	Government	UKGov_SDA_designations
24	*France 24 / RFI*	Ilya Gambashidze: Simple soldier of disinformation or king of Russia's trolls?	2024-02-28	Journalism	France24_Gambashidze_2024-02
25	*German Federal Foreign Office (Auswärtiges Amt)*	Germany Targeted by the Pro-Russian Disinformation Campaign "Doppelgänger" — Technical Report on an Analysis by the Federal Foreign Office	2024-06-05	Government	German_FO_technical_2024
26	*Global Affairs Canada (Government of Canada)*	G7 Rapid Response Mechanism (RRM) statement on Russian Influence Campaign	2025-01-17	Government	G7RRM_Canada_2025-01
27	*Gnida Project*	Recent Social Design Agency leak points to	2026-05-30	Research / OSINT	Gnida_SDA_Storm1516_2025

No.	Publisher	Title	Date	Type	Identifier
		Storm-1516 coordination with the Presidential Administration			
28	***Google Threat Analysis Group (TAG) / Trust & Safety***	TAG Bulletin: Q1 2025	2025-05-15	Platform	GoogleTAG_Bulletin_Q1_2025
29	***Graphika***	Summit Old, Summit New: Russia-Linked Actors Leverage New and Old Tactics in Influence Operations Targeting Online Conversations About NATO Summit	2023-08-21	Research / OSINT	Graphika_Summit
30	***Hive Pro (HiveForce Labs)***	Unmasking Doppelgänger: Russia's Disinformation Campaign Revealed	2024-02-27	Research / OSINT	HiveForce_UnmaskingDoppelganger
31	***HM Treasury / Office of Financial Sanctions Implementation (OFSI)***	OFSI Designation: Social Design Agency (Unique ID RUS2292)	2024-10-28	Government	RUS2292_SDA_2024
32	***IFEX (republishing European Federation of Journalists / MFRR joint statement)***	Russian "Doppelgänger" campaign spoofs media websites to spread disinformation	2024-09-30	Self-report	SR_IFEX
33	***Institute for Strategic Dialogue (ISD)***	Russian influence operation Doppelganger linked to fringe advertising company	2024-06-13	Research / OSINT	ISD_ArgonLabs

No.	Publisher	Title	Date	Type	Identifier
34	***Institute for Strategic Dialogue (ISD)***	Investigation | How a pro-Kremlin ad campaign used the Israel-Hamas conflict to spread propaganda in France and Germany	2024-10-11	Research / OSINT	ISD_IsraelHamas_2024-10
35	***Institute of Mass Information (IMI)***	Criminals are conducting a fake survey under the guise of RBC-Ukraine	2022-07-11	Self-report	SR_RBC_imi
36	***Institute of Mass Information (IMI)***	Scammers pretending to be the "Obozrevatel" are promoting fake posts saying that "the allies are betraying us"	2023-04-11	Self-report	SR_Obozrevatel
37	***Intrinsec***	Doppelgänger: New disinformation campaigns spreading on social media through Russian networks	2025-02-27	Research / OSINT	Intrinsec_Doppelganger
38	***Lund University & Swedish Psychological Defence Agency (MPF Report Series 8/2025)***	Beyond Operation Doppelgänger: A Capability Assessment of the Social Design Agency	2025-05-12	Academic	Lund_BeyondDoppelganger
39	***Meduza***	'Latching onto successful projects': Leaked documents suggest Kremlin spin doctors are presenting popular movies and TV to Putin as propaganda wins	2024-02-26	Journalism	Meduza_KremlinLeaks_2024-02
40	***Meta***	Adversarial Threat	2024-	Platform	Meta_ATR_2024-Q1

No.	Publisher	Title	Date	Type	Identifier
		Report, First Quarter 2024	05-01		
41	***Meta***	Q3 2024 Adversarial Threat Report	2024-12-01	Platform	Meta_ATR_2024-Q3
42	***Meta Platforms, Inc.***	Q4 2022 Adversarial Threat Report	2023-02-01	Platform	Meta_ATR_2022-Q4
43	***Microsoft On the Issues / Microsoft Threat Analysis Center (MTAC)***	Russian election interference efforts focus on the Harris-Walz campaign	2024-09-17	Research / OSINT	MTAC_HarrisWalz
44	***Microsoft Threat Analysis Center (MTAC)***	Nation-states engage in US-focused influence operations ahead of US presidential election	2024-04-17	Research / OSINT	MTAC_Elections_2024-04
45	***Microsoft Threat Analysis Center (MTAC)***	Russia leverages cyber proxies and Volga Flood assets in expansive influence efforts	2024-09-17	Research / OSINT	MTAC_ElectionReport4
46	***Microsoft Threat Analysis Center (MTAC)***	Russia, Iran, and China continue influence campaigns in final weeks before Election Day 2024	2024-10-23	Research / OSINT	MTAC_ElectionReport5
47	***OCCRP (with Delfi Estonia)***	Leaked Documents Reveal Russian 'Cognitive Strikes' Against the West — Including Islamophobic 'Pig Head' Attacks in Paris	2026-05-24	Leak-derived	OCCRP_SDAleak_2026
48	***OpenAI***	Influence and cyber operations: an update, October 2024	2024-10	Research / OSINT	OpenAI_2024-10

No.	Publisher	Title	Date	Type	Identifier
49	***OpenAI***	Disrupting malicious uses of our models: an update, February 2025	2025-02	Research / OSINT	OpenAI_2025-02
50	***OpenAI***	Disrupting malicious uses of AI: June 2025	2025-06	Research / OSINT	OpenAI_2025-06
51	***OpenAI***	Disrupting malicious uses of our models: an update, October 2025	2025-10	Research / OSINT	OpenAI_2025-10
52	***OpenAI***	Disrupting malicious uses of AI	2026-02	Research / OSINT	OpenAI_2026-02
53	***Qurium Media Foundation***	Under the hood of a Doppelgänger	2022-09-27	Research / OSINT	Qurium_UnderHood
54	***Qurium Media Foundation***	How Russia uses EU companies for propaganda - Exposing The Evil Empire of Doppelganger Disinformation	2024-07-11	Research / OSINT	Qurium_EvilEmpire
55	***Qurium Media Foundation***	When Kehr meets VexTrio	2024-11-13	Research / OSINT	Qurium_KehrVexTrio_2024
56	***RBC-Ukraine (RBC.UA / LLC UBT MEDIA)***	Увага! Зловмисники проводять фейкове опитування від імені РБК-Україна	2022-07-11	Self-report	SR_RBC_primary
57	***Recorded Future (Insikt Group)***	Russia-Linked CopyCop Uses LLMs to Weaponize Influence Content at Scale	2024-05-09	Research / OSINT	GT-0109_CopyCop
58	***Recorded Future (Insikt Group)***	“Operation Undercut” Shows Multifaceted Nature of SDA’s Influence Operations	2024-11-26	Research / OSINT	RecordedFuture_OperationUndercut

No.	Publisher	Title	Date	Type	Identifier
59	***Recorded Future (Insikt Group)***	Stimmen aus Moskau: Russian Influence Operations Target German Elections	2025-02-13	Research / OSINT	RecordedFuture_StimmenAusMoskau
60	***Recorded Future (Insikt Group)***	CopyCop Deepens Its Playbook with New Websites and Targets	2025-09-17	Research / OSINT	RecordedFuture_cta-ru-2025-0917
61	***Recorded Future — Insikt Group***	Obfuscation and AI Content in the Russian Influence Network "Doppelgänger" Signals Evolving Tactics	2023-12-05	Research / OSINT	RecordedFuture_Obfuscation_AI_2023
62	***Reset Tech***	So We Meet Again, Botiful: Latent Fake Ad Network Persists on Meta Two Years Since Doppelganger	2024-11-12	Research / OSINT	ResetTech_Botiful_2024
63	***Reset Tech***	Verified Disinformation: How X Profits From the Rise of a Pro-Kremlin Network	2024-12-01	Research / OSINT	ResetTech_VerifiedDisinformation
64	***Reuters (investigations)***	Russia-linked propaganda campaign pushes to undercut German support for Ukraine	2025-02-18	Journalism	GT-0392_undercut_German
65	***RFE/RL (Schemes, investigative unit of RFE/RL's Ukrainian Service)***	Leaked Records Detail Vast Russian Influence Campaign Targeting Ukraine, EU	2024-09-17	Journalism	RFERL_Schemes_SDAleak_2024

No.	Publisher	Title	Date	Type	Identifier
66	*Science Feedback & Alliance4Europe (with Logically Facts, ASD/GMF, Prose Intelligence, DRI, et al.)*	Sanctioned but Thriving: How Online Platforms Fail To Address the Widespread Presence of Entities Under EU Sanction	2024-12-03	Academic	SanctionedButThriving
67	*Sekoia Threat Detection & Research (TDR)*	Master of Puppets: Uncovering the DoppelGänger pro-Russian influence campaign	2024-05-14	Research / OSINT	Sekoia_MasterOfPuppets
68	*SentinelLABS (SentinelOne) & ClearSky Cyber Security*	Doppelgänger | Russia-Aligned Influence Operation Targets Germany	2024-02-01	Research / OSINT	SentinelLabs_TargetsGermany
69	*t-online*	Fake-Seiten und Bots: Diese Russen arbeiten für Kampagne gegen den Westen	2024-03-27	Journalism	TOnline_initial_exposure_2022
70	*The Forward*	The Forward is targeted by Russian disinformation campaign around Israel	2024-06-27	Self-report	SR_Forward
71	*The Insider*	German investigative journalists confirm Doppelgaenger bot network's digital trails lead to Russian Ministry of Defense	2024-11-18	Journalism	TheInsider_MoD_2024-11
72	*U.S. Department of Justice / Federal Bureau of Investigation (Eastern District of Pennsylvania)*	Affidavit in Support of Seizure Warrant (In re Certain Domains), Case No. 24-mj-1395 (E.D. Pa.)	2024-09-04	Government	DOJ_affidavit_full

No.	Publisher	Title	Date	Type	Identifier
73	***U.S. Department of Justice, Office of Public Affairs***	Justice Department Disrupts Covert Russian Government-Sponsored Foreign Malign Influence Operation Targeting Audiences in the United States and Elsewhere	2024-09-04	Government	DOJ_seizure_pressrelease
74	***U.S. Department of the Treasury (OFAC)***	United States, Australia, and United Kingdom Sanction Russian Cybercrime Infrastructure Supporting Ransomware	2025-11-19	Government	OFAC_MediaLand_Hypercore_2025
75	***U.S. Department of the Treasury, Office of Foreign Assets Control (OFAC)***	Treasury Sanctions Actors Supporting Kremlin-Directed Malign Influence Efforts	2024-03-20	Government	OFAC_SDA_Structura_2024
76	***U.S. Department of the Treasury, Office of Foreign Assets Control (OFAC)***	Treasury Sanctions Global Bulletproof Hosting Service Enabling Cybercriminals and Technology Theft	2025-07-01	Government	OFAC_Aeza_2025
77	***UK Foreign, Commonwealth & Development Office (FCDO)***	UK sanctions Putin's interference actors	2024-10-28	Government	UKFCDO_OFSI_sanctions_2024
78	***United States Department of State***	Imposing Sanctions on Actors Supporting Kremlin-Directed Disinformation Efforts	2024-03-20	Government	USState_KremlinDisinfo_2024-03
79	***US Department of the Treasury (OFAC)***	Treasury Takes Action as Part of a U.S. Government	2024-09-04	Government	OFAC_JY2559_2024-09

No.	Publisher	Title	Date	Type	Identifier
		Response to Russia's Foreign Malign Influence Operations			
80	***US Department of the Treasury, Office of Foreign Assets Control (OFAC)***	Notice of OFAC Sanctions Actions (SDN designations of 20 March 2024)	2024-03-25	Government	FederalRegister_OFAC_Mar2024
81	***VIGINUM (Service de vigilance et de protection contre les ingerences numeriques etrangeres), SGDSN, France***	RRN: A complex and persistent information manipulation campaign	2023-07-19	Government	VIGINUM_RRN_2023
82	***VIGINUM (SGDSN)***	War in Ukraine: Three years of Russian information operations	2025-02-01	Government	VIGINUM_ThreeYears
83	***VIGINUM (SGDSN, France)***	Portal Kombat: A structured and coordinated pro-Russian propaganda network (Part Two)	2024-02	Government	VIGINUM_PortalKombat_P2
84	***VIGINUM (SGDSN, France)***	Portal Kombat: A structured and coordinated pro-Russian propaganda network	2024-02-12	Government	VIGINUM_PortalKombat_P1
85	***VIGINUM (SGDSN, France)***	Portal Kombat — Expansion of the pro-Russian propaganda network: new domain names (Part Three, Technical report)	2024-04-28	Government	VIGINUM_PortalKombat_P3
86	***VIGINUM (SGDSN,***	Information	2026-	Government	VIGINUM_IMS_2026

No.	Publisher	Title	Date	Type	Identifier
	France)	Manipulation Set (IMS) — concept definition report	01-22		
87	*VSquare / Delfi Estonia / FRONTSTORY.PL*	Leaked Files from Putin's Troll Factory: How Russia Manipulated European Elections	2024-09-16	Journalism	VSquare_LeakedFiles_2024-09
88	*VSquare.org (republication of OCCRP / Delfi Estonia investigation)*	Leaked Russian Docs Reveal 'Cognitive Strikes' on the West – Including 'Pig Head' Attacks in Paris	2026-05-27	Leak-derived	VSquare_SDAleak_2026

Annex B - The 33 binding measures

Every binding measure addressing an element of SDA's operation, one row per act and in date order, with standing and undated measures last. Where several registry records describe the same act they are consolidated, and the records are listed so that each act can be traced. Three acts touch more than one element. Where the records of an act disagree on its outcome, the best-evidenced assessment is given and the act is marked as contested.

Act	Date	Authority	Measure	Element	Recorded outcome	Registry records
A01	2022-09-27	*Meta*	Takedown of the Doppelgänger network: accounts, pages, ads and 60+ impersonating sites	Clone domains; Social accounts and bots; Ad-buying and payment	degraded	D096, D009, D011
A26	2022-11	*WIPO*	Transfer of	Clone	not	D104

Act	Date	Authority	Measure	Element	Recorded outcome	Registry records
		(UDRP)	sueddeutsche[.]me to Süddeutsche Zeitung	domains	assessed	
A27	2023-07	***WIPO (UDRP)***	Transfer of lemonde[.]ltd to Le Monde	Clone domains	not assessed	D115
A02	2023-07-28	***Council of the EU***	EU listing of SDA, Structura and individuals	The organisation itself	degraded (contested)	D050, D053, D054, D055, D062, D068, D061
A28	2023-10	***WIPO (UDRP)***	Order that NameCheap return gouv[.]fm to the French foreign ministry	Clone domains	not assessed	D119
A08	2023-11	***Social media platforms***	Suspension of Doppelgänger CIB accounts	Social accounts and bots	adapted to continue (contested)	D121, D010
A10	2024	***Meta***	Blocking of Doppelgänger domains, ads and links across apps	Clone domains; Ad-buying and payment	degraded	D014, D099, D110
A11	2024	***Netshield***	Deletion of hosted Doppelgänger domains	Clone domains	not assessed	D127
A09	2024-01-31	***X***	Suspension of exposed inauthentic accounts	Social accounts and bots	adapted to continue	D140
A29	2024-02	***AFNIC***	Return of leparisien[.]re to Le Parisien	Clone domains	not assessed	D147
A03	2024-03-20	***US Treasury (OFAC)***	US designation of SDA, Structura and their founders	The organisation itself	degraded (contested)	D033, D040, D041, D150, D151, D153, D155,

Act	Date	Authority	Measure	Element	Recorded outcome	Registry records
						D025, D026
A12	2024-07-01	***Hetzner***	Closure of Doppelgänger customer accounts	Hosting and IP	not assessable	D161
A13	2024-07-24	***Aeza International***	Closure of the Lethost and Sunhost bulletproof subsidiaries	Hosting and IP	adapted to continue	D165
A14	2024-07-31	***Datapacket***	Withdrawal of upstream connectivity from Aeza	Hosting and IP	degraded	D166
A04	2024-09-04	***US DOJ / FBI***	Seizure of 32 Doppelgänger domains, on the FBI affidavit	Clone domains	degraded (contested)	D001, D003, D004, D005, D007, D132, D168, D006, D169, D171
A16	2024-10	***Ukrainian provider***	Termination of the Russian customer	Hosting and IP	degraded	D174
A17	2024-10	***X***	Suspension of accounts linked to the new domains	Social accounts and bots	adapted to continue	D141
A15	2024-10-01	***Kehr operator***	Blocking of Doppelgänger and closure of the cloaking service	Hosting and IP	degraded	D176
A05	2024-10-28	***UK FCDO / OFSI***	UK designation of SDA, Structura, ANO Dialog and individuals	The organisation itself	degraded	D036, D073, D074, D081, D175
A19	2024-11	***European telecommunication***	Cut-off of SDA and blocking of RRN	Hosting and IP	degraded	D178

Act	Date	Authority	Measure	Element	Recorded outcome	Registry records
		s carriers				
A20	2024-11	***German company***	Takedown of campaign sites including RRN	Clone domains	stopped	D179
A32	2024-11	***Apliteni (Keitaro)***	Revocation of the Keitaro licences used for cloaking, reported 15 Nov 2024	Hosting and IP	not assessed	D146
A18	2024-11-07	***1API GmbH***	Blocking of 28 Doppelgänger domains	Clone domains	stopped	D180
A06	2024-12-16	***Council of the EU***	EU hybrid-threats listing including Doppelgänger-associated entities	The organisation itself	not assessed	D057, D049
A33	2025-02	***Google***	Delisting of 47 Doppelgänger domains from Google News and Discover	Clone domains	not assessed	D186
A21	2025-07-01	***US Treasury (OFAC)***	US designation of Aeza Group, affiliates and four individuals	Hosting and IP	no change	D156, D037
A22	2025-11	***UK FCDO / OFSI***	UK designation of Aeza Group for hosting SDA campaigns	Hosting and IP	not assessed	D078
A23	2025-11-19	***US Treasury (OFAC)***	US designation of Aeza successor companies	Hosting and IP	not assessed	D039
A07	2026-05	***UK FCDO / OFSI***	UK designation of SDA and 49 staff	The organisation itself	not assessed	D072, D080
A24	standing	***X***	Blacklisting of Doppelgänger domains	Clone domains	adapted to continue	D201
A25	standing	Meta	Blocking of fake accounts and advertiser pages	Social accounts and	no change (contested)	D101, D013,

Act	Date	Authority	Measure	Element	Recorded outcome	Registry records
				bots; Ad-buying and payment		D016, D107
A31	undated	OpenAI	Termination of the operation's accounts used to produce content	Fake-brand content	degraded	D203
A34	standing	X	Suspension of accounts and removal of verification badges in a Doppelgänger network	Social accounts and bots	no change	D142