ISBN 978-9934-619-68-7

ATTRIBUTING RUSSIAN INFORMATION INFLUENCE OPERATIONS

Testing the Information Influence Attribution Framework with real-world case studies

A JOINT REPORT FROM
THE UKRAINIAN CENTRE FOR STRATEGIC COMMUNICATIONS AND THE NATO STRATEGIC COMMUNICATIONS CENTRE OF EXCELLENCE

CENTRE FOR STRATEGIC COMMUNICATIONS

ISBN: 978-9934-619-68-7

Authors: Victoria Smith, Dr James Pamment, Sofiia Dikhtiarenko, Ben Heap,
Darejan Tsurtsumia, Adam Maunder.
Project Manager: Ben Heap.
With thanks to: Maria Sahaidak, Iryna Subota, Casper Haanappel.
Layout: Una Grants

Contact us at Ukraine@stratcomcoe.org

January 2026

NATO STRATCOM COE
11b Kalnciema iela,
Riga, LV1048, Latvia
stratcomcoe.org
@stratcomcoe

ATTRIBUTING RUSSIAN INFORMATION INFLUENCE OPERATIONS

Testing the Information Influence Attribution Framework with real-world case studies

Contents

Foreword

This joint report provides a structured, multi-source framework to attribute responsibility for Information Influence Operations, specifically focusing on Russian campaigns targeting Ukraine and its neighbours. Our analysis refines existing attribution processes, aiming to produce conclusions that are credible, actionable, and transparent.

Attribution here is not only about identifying responsible parties; it seeks to empower Ukraine and its partners to challenge hostile narratives, expose sources of manipulation, and undermine adversary legitimacy. By assembling robust evidence, decision-makers can hold malign actors accountable and justify proportional responses, ranging from public exposure to legal and diplomatic action.

The report is anchored in practical case studies drawn from recent Russian influence operations, analysed with technical, behavioural, and contextual evidence. Special attention is given to audiences most impacted, including Ukrainian civilians, regional allies, and European pro-Kremlin groups.

The scope covers assessment tools, standards, and frameworks, including the DISARM methodology and the contributions of governments, civil society, and technology platforms. The goal is clear: equip organisations to systematically detect and counter foreign influence through rigorous attribution, supporting legal reforms and policy initiatives such as the Digital Services Act, and thus strengthening democratic resilience.

Ben Heap, NATO Strategic Communications Centre of Excellence.
Mykola Balaban, Centre for Strategic Communications, Ukraine.

Executive Summary

This report examines Russian Information Influence Operations targeting audiences in Ukraine and neighbouring regions, including Ukrainian civilians and defence forces, civilians in nearby states, and European pro-Kremlin groups. The analysis draws primarily on data from the Ukrainian Centre for Strategic Communications (CSC), complemented by recent government and civil society reporting.

The aim is to test and refine the Information Influence Attribution Framework (IIAF) by applying it to real-world Russian campaigns, in a context where EU sanctions on Russian state media, the Foreign Information Manipulation and Interference (FIMI) policy framework, and the Digital Services Act (DSA) are raising evidential standards for attribution. As Information Influence Operations increasingly involve both governmental and civil-society actors, the report focuses on clarifying practical evidential thresholds and confidence levels that can withstand prospective legal and regulatory scrutiny.

Chapter 2 introduces the attribution framework used throughout the report. It categorises evidence as technical (digital traces and infrastructure metadata), behavioural (tactics, techniques, and procedures), and contextual (narratives, timing, and political environment), and adds a legal-ethical assessment to weigh proportionality, data protection, and geopolitical considerations. The chapter explains how these evidence types can be drawn from open, proprietary, and classified sources, and how confidence intervals and a spectrum of state responsibility help express uncertainty and degrees of state involvement.

Chapter 3 applies this framework to Russian operations related to Ukraine. Case studies cover RT and Sputnik's post-sanctions infrastructure workarounds, coordinated Telegram networks amplifying pro-Kremlin narratives, cross-posting and source-laundering around a fabricated clash between Georgian and Ukrainian soldiers, and a DISARM-mapped campaign pushing false claims that Poland seeks to annex western Ukraine. A corruption-focused operation is analysed in depth, showing how converging technical, behavioural, and contextual indicators, together with state-responsibility and confidence scales, support a high-confidence assessment that the campaign is state-shaped to state-coordinated by the Russian Federation.

Chapter 4 distils lessons learned for practitioners. It highlights the structural limits of open-source and partial platform data, the necessity of documenting confidence ranges and evidential gaps transparently, and the value of standardised language (e.g. state-encouraged, state-shaped, state-coordinated) in public-facing attributions. Recommendations include the improvement of mechanisms for secure access to proprietary and classified data, strengthening cooperation among governments, platforms, and civil society, and further systematising network, TTP, and narrative analysis so that future attributions are more robust against legal challenge under instruments such as the DSA and related EU counter-FIMI measures.

Introduction

This research focuses on Russian Information Influence Operations and their efforts to shape perceptions in Ukraine and its regional neighbourhood.[1] Their targeted audiences include Ukrainian civilians, Ukrainian Armed Forces, people in neighbouring EU countries providing humanitarian and military support to Ukraine, as well as European audiences sympathetic to, or uncertain about, Kremlin-aligned narratives.

The focus of this report is on evaluating and refining the Information Influence Attribution Framework (IIAF) by applying it to real-world case studies of Russian influence campaigns. The analysis takes place within the broader policy context of the European Union's ongoing counter-FIMI initiatives: the Foreign Information Manipulation and Interference (FIMI) policy framework, the March 2025 introduction of the FIMI Exposure Matrix, and the continued enforcement of the Digital Services Act (DSA) since February 2024. Collectively, these policy measures illustrate how Western democracies have intensified their collective efforts to detect, expose, and counteract foreign information manipulation.

Attributing Information Influence Operations implies meeting an evidential threshold. This threshold will increasingly be tested because of policy innovations. In other words, we are increasingly likely to see 'lawfare' [2] conducted by organisations connected to the Russian Federation's influence apparatus [3] targeting, for example, regulatory decisions and de-platforming actions through litigation.

Given that attribution analyses are conducted through collaborative efforts between national authorities, private-sector investigators, and civil society (including non-governmental organizations, journalists, think tanks, and research institutions), it is critical to maintain transparent methodological standards. This need is heightened by the reduction of accessible social media and messaging platform data for external researchers following API restrictions imposed in 2024 and 2025, which have hindered open-source investigative capabilities.[4]

The Information Influence Attribution Framework (IIAF) outlines a process of analysing technical, behavioural, and contextual evidence, together with a legal and ethical assessment in order to establish the source.[5] The framework distinguishes:

- **Technical evidence**. Observable traces an adversary leaves behind, such as digital signals, telemetry, financial, or other traceable physical evidence.
- **Behavioural evidence**. Supported by knowledge of the methods by which different adversaries carry out their work (often termed Tactics, Techniques and Procedures or TTPs).
- **Contextual evidence**. Analysis of the content of an influence operation, the socio-political context it seeks to influence, and motivations of the adversary.
- **Legal and ethical assessment**. An assessment of whether assigning blame is proportionate, and whether it sets into motion considerations relating to political or commercial fallout, treaties, or litigation.

By applying the IIAF to Russian Information Influence Operations around Ukraine, this research clarifies practical evidential thresholds for credible attribution and identifies best practices for state, private-sector, and civil society practitioners.

Key terms

Influence Operations A coordinated set of activities to influence the perceptions or behaviour of a target audience and achieve a specific goal to the benefit of the influencer.

FIMI Foreign Information Manipulation and Interference (FIMI) describes a mostly non-illegal pattern of behaviour that threatens or has the potential to negatively impact values, procedures and political processes. Such activity is manipulative in character, conducted in an intentional and coordinated manner, by state or non-state actors, including their proxies inside and outside of their own territory.

Digital Services Act (DSA) The DSA regulates online intermediaries and platforms such as marketplaces, social networks, content-sharing platforms, app stores, and online travel and accommodation platforms. Its main goals are to prevent illegal activities online, mitigate systemic risks, and increase transparency and accountability in content moderation.

TTPs 'Tactics, Techniques, and Procedures' describe the patterns of behaviour used by threat actors to manipulate the information environment with the intention to deceive. DISARM incorporates 'Tactics,' which are the operational goals that threat actors are trying to accomplish and 'Techniques,' which are the actions through which they try to accomplish them. 'Procedures' describe the specific combination of techniques across multiple tactics (or stages of an attack) that indicate intent and may be unique for different threat actors.

SIGINT SIGnals INTelligence is collected from electronic signals and systems, usually by governments.

HUMINT HUMan INTelligence is collected directly from people.

OSINT OpenSource INTelligence is the process of collecting and analysing intelligence in the public domain to answer specific intelligence-led questions.

API An application programming interface is a software intermediary that allows two different applications to talk to each other. APIs are a way of extracting and sharing data within and across organizations.

OpenCTI Open Cyber Threat Intelligence is an open-source platform for structuring, storing, and analysing information on cyber and disinformation threats. Data structures and relationships are based on the STIX standard and can be analysed using network graphs. A key feature of OpenCTI is the ability to connect the dots connecting new and existing threat information.

STIX The Structured Threat Information Expression (STIX) language is a data format used to encode and exchange cyber threat intelligence (CTI). It can also be used to share information on FIMI incidents, by breaking them down into their different constitutive elements. STIX objects are used to represent threat actors, the techniques they use, the narratives they propagate, the channels they use etc. STIX objects are stored and analysed in threat intelligence platforms such as OpenCTI

DAD-CDM Common Data Model for Defending Against Disinformation (DAD-CDM). This is an open project, available to any individual or organisation who wishes to help in the creation of a common data model for characterising and responding to threats involving manipulation of the information environment. The DISARM Foundation initiated this project with OASIS Open, which aims to model disinformation threats and responses by building upon the STIX standard with disinformation-specific objects, properties, and relationships.

Understanding attribution

Attribution is about ***perceptions of cause***. When we attribute something, we make a claim about its origins. The art world depends on accurate attribution of authorship to distinguish between the genuine and the fake. In relation to Information Influence Operations (IIOs), attribution means identifying the threat actor responsible. This can be challenging for several reasons, such as a lack of access to relevant data, the difficulties of uncovering tracks that have been deliberately hidden, and of distinguishing the originator of a campaign from others who disseminate similar or identical content.

The field of information influence research is further complicated by a lack of clarity about who is permitted to lie or mislead in public debate. In most liberal democracies, citizens and people living in a country have different rights to, for example, a foreign news agency. In the language of the Digital Services Act,[6] the European Union refers to Foreign Information Manipulation and Interference (FIMI), and it is the 'foreign' part of FIMI that is often the key component that motivates further investigation of questionable content, but to ascertain 'foreignness', and then to develop a response, we need attribution.[7]

Attribution debates take their inspiration from the cyber security field, where the ***technical*** evidence to support attribution of cyberattacks is often more clearcut, since threat actors use specific tools and methods that are hallmarks of their approach to system intrusion. This is not dissimilar to artists whose work can be identified by attention to certain brushstrokes or the layering of paint.

IIOs aren't as straightforward to attribute since the possible 'fingerprints' of a campaign tend to be analysed in patterns of ***behaviour*** (how the campaign is conducted) and patterns of ***discourse*** (the content of the campaign, which includes political context). Many of these aspects are either highly generic (i.e., the behaviours and content do not stand out from the crowd), or there is simply no way of distinguishing a source from the available data.

The results of an investigation can be strengthened by intelligence collected through signals intelligence (SIGINT) or human intelligence (HUMINT), or through the proprietary data collected through digital platform backends (often referred to as telemetry) or by searching through databases for example that connect IP addresses, telephone numbers, and email addresses to companies and individuals.

In cases that do not meet the threshold for law enforcement or intelligence agencies to get involved, investigators of influence operations must accept that they are working with limited data, often involving major gaps in available information.

However, it is still possible to build **a credible and compelling case** that can demonstrate a connection to a hostile actor, even if it can't identify which specific organisations or individuals are likely to be responsible.

The Attribution Framework

According to the NATO StratCom COE and Hybrid COE framework (2022),[8] attribution can be theorised using three types of evidence: **technical**, **behavioural,** and **contextual**, supported by a **legal and ethical assessment**. This structure is retained and further operationalised in the later ADAC.io IIO Attribution Framework.[9]

Technical evidence focuses on the trail of signals generated by illicit activities, such as IP addresses.

Behavioural evidence focuses on manipulative activities and techniques, including Tactics, Techniques and Procedures (TTPs).

Contextual evidence examines content and political elements such as messaging and narratives.

Finally, the **legal and ethical assessment** weighs up crucial questions of proportionality, data protection, and geopolitical strategy related to using these different kinds of evidence.

Each category of evidence can be subdivided into types of data source. Evidence can be collected through **open sources** (e.g. through research, open access APIs, and OSINT), **proprietary sources** in which the data has commercial ownership (e.g. social media platform backends and API's, private sector intelligence), and through **classified intelligence** (e.g. SIGINT and HUMINT).

Variations in attribution accuracy

- Actors based in X country (the attribution only goes so far as to identify the territory from which the operation was staged)
- Individuals associated with Y organisation (the attribution identifies an organisation such as a foreign military but is unable to ascertain whether the operation was conducted as a matter of governmental policy)
- Y organisation associated with or acting on behalf of X country (the attribution establishes ties between the operation and a government, for example through procurement contracts)
- Z individuals working for Y organisation on behalf of country X (the attribution was able to reveal who worked on the operation and under which authority)

	Technical evidence	Behavioural evidence	Contextual evidence	Legal and ethical assessment
Open source	Web domain ownership, IP addresses, economic ties	Account activity, page activity, posting/ cross-posting, sharing, follows, network	Media content, discourse and narratives, linguistics, political context, cui bono	Risk of litigation; research ethics; personal risk of becoming a target
Proprietary source	Data collected by platform backend Data collected by platform backend	As above, with more extensive platform data	As above and data on previous takedowns with suspected linksAs above and data on previous takedowns with suspected links	Protecting political and commercial interests; data protection
Classified source	SIGINT; proprietary source data acquired by warrant	As above and SIGINT, HUMINT	As above and classified geo-political assessmentsAs above and classified geo-political assessments	Actor-specific strategy; protecting political interests; data protection

Matrix of evidence and data sources[10]

Organisations conducting analysis rarely have access to all these types of data; open source and proprietary data are usually the main evidence types referred to in public attributions.

The purpose of the Information Influence Attribution Framework is twofold. First, it helps to demonstrate that an attribution consists of several assessments that in combination help to build a credible picture. There may be a considerable weight of evidence in some categories, and next to none in the others. It may be possible to find similar technical evidence through both open and classified sources. Even then, legal and ethical assessments may lead to a decision not to attribute, for example to protect sources when secret intelligence is a core factor in the attribution.

Second, the matrix provides a means of communicating and sharing high-level, non-specific data on the main factors that constitute an attribution. An actor could use the matrix to highlight that open-source contextual data is the principal evidence that has been used to justify a decision, thereby offering a little more nuance to stakeholders about the basis for the attribution.

Technical evidence

Technical analysis provides a structured approach for identifying and interpreting digital artefacts, such as infrastructure metadata or platform-level signals, that reveal how Information Influence Operations (IIOs) are built, executed, and sustained.

These artefacts, **technical evidence,** can be processed using analytical tools to detect anomalies or consistencies that indicate coordination, centralised control, or deliberate obfuscation.

Observable infrastructure traces include telemetry, metadata and hosting or account infrastructure, such as domains, IP addresses, account statistics, and transaction records that can be observed and measured.

Technical evidence can reveal patterns where they shouldn't be (e.g., synchronised cross-posting across multiple accounts) or highlight the absence of patterns where consistency would be expected (e.g., accounts lacking metadata typically present on similar platforms).

It is among the clearest entry points for detection and attribution, offering a relatively objective foundation for analysis that can be assessed for anomalies or consistencies. However, it must be supported with behavioural and contextual evidence to build high-confidence attribution.

Technical analysis works across three evidence categories:

Digital infrastructure the underlying technical components of an operation, such as domain names, IP addresses, hosting services, DNS records, SSL certificates, and other technical artefacts;

Platforms and networks technical metadata related to how content circulates across social platforms. This includes account metrics such as subscriber counts, posting frequency, and engagement rates, as well as repost timing, cross-channel mentions, and shared follower patterns;

Financial and commercial signals help uncover the funding and monetisation of operations through indicators such as blockchain activity and advertising infrastructure.

The following sections explain in more details how technical analysis is conducted in each of these categories.

Digital infrastructure

Digital infrastructure refers to the tech-
Digital infrastructure refers to the technical foundations that support and distribute content online. This includes hosting locations (IP addresses), registry details, naming conventions (subdomains and DNS entries), and cryptographic metadata (SSL certificates). These elements often reveal who controls a given online asset and whether that infrastructure is reused across multiple operations. Repetition, reuse, or close registration timing may indicate central coordination or mass deployment.

An assessment of digital infrastructure begins with the identification of relevant assets, such as websites, accounts, or servers linked to suspicious activity. Tools like WHOIS lookups, passive DNS databases, and historical domain records can then be used to uncover metadata about ownership, configuration, and history.

A WHOIS lookup of the domain *fondfbr.ru* (see screenshot) illustrates the type of metadata analysts can access during infrastructure analysis.[11] The domain is registered to the Foundation for Combating Repression (Фонд борьбы с репрессиями), a Russian organisation reported to be affiliated with Yevgeny Prigozhin's network.

The WHOIS results reveal details such as domain creation date, registrar (e.g. REG. RU), nameservers, and sometimes anonymised or visible registrant data. These details help analysts trace relationships between domains, identify potentially coordinated deployments, or spot attempts at obfuscation.

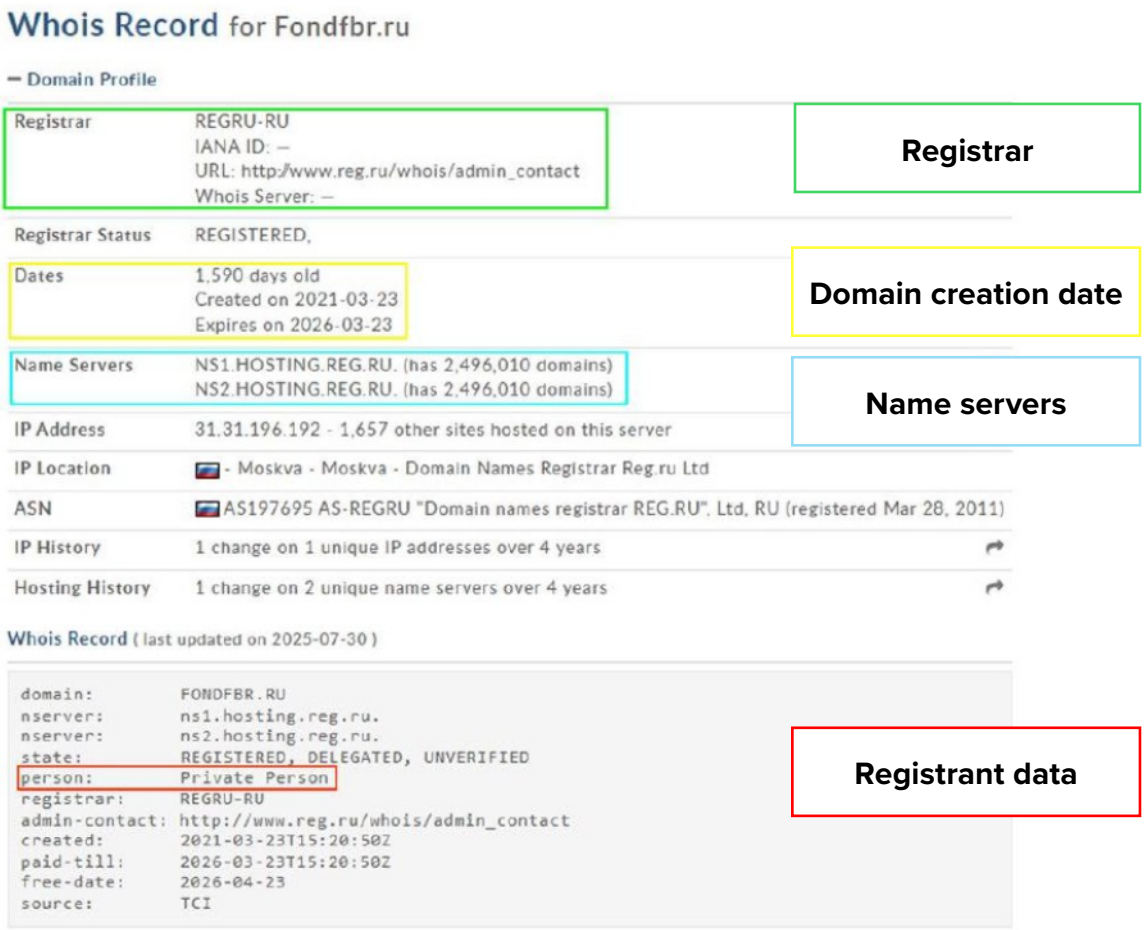

WHOIS lookup results for fondfbr.ru

Analysts then identify and compare patterns; for instance, shared IP addresses between unrelated sites may indicate central hosting; identical registrars or name server configurations can imply shared management; and closely timed domain registrations help construct deployment timelines. These indicators are measured not in isolation but across clusters – anomalies, repetitions, and convergences are what point to campaign-level control. Analysts can also track elements like anonymised registrant entries or extended registration periods, which may signal attempts to obfuscate or sustain long-term operations.

Technical analysis of digital infrastructure can produce traceable, machine-readable evidence. Single indicators can be coincidental but consistent patterns across multiple sources might confirm operational reuse and establish the structural backbone of an IIO. This helps attribution by narrowing down likely operators or exposing shared infrastructure behind seemingly unrelated fronts.

Platforms and networks

This category focuses on the technical metadata collected from social media and messaging platforms to examine the distribution architecture of influence operations. It includes statistical analysis of platform usage, account behaviours, and dissemination patterns that can indicate central coordination or automation.

Assessment begins with the collection of **platform data** (also known as metadata) which includes details about account behaviour and interactions. Tools such as *TGStat* (see screenshot below), a web platform for retrieving public data from Telegram channels, can gather data such as subscriber counts, growth trends, average posting frequency and Engagement Rate by Reach (ERR). This data forms a baseline for understanding channel ownership, activity and content distribution.

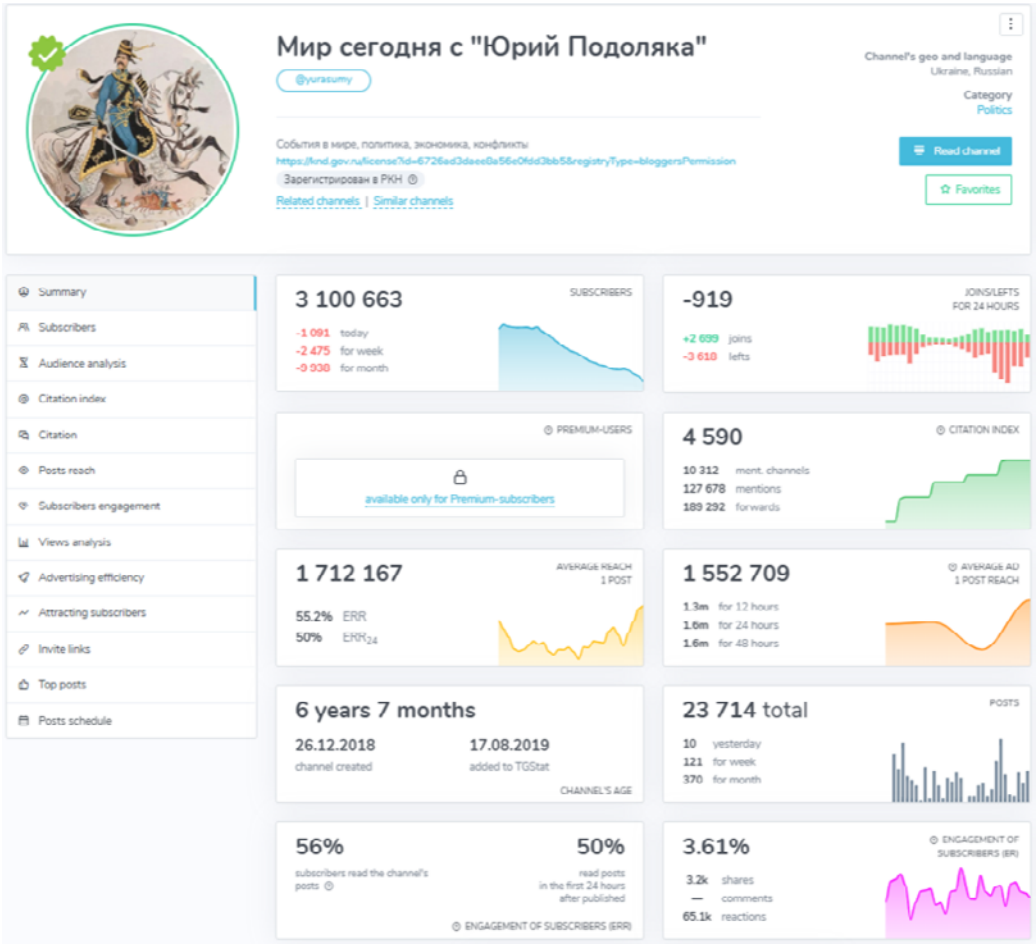

TGStat profile for @yurasumy 'The World Today with Yuriy Podolyaka' [12]

Next is **network mapping** – tracking how messages are forwarded, which accounts mention each other and where content originates using citation graphs and reposting intervals. TGStat tool includes detailed forwarding and citation graphs, which illustrate how the specific channel's posts are disseminated across the Telegram ecosystem.

The screenshot below displays the incoming and outcoming mentions for *@yurasumy* ('The World Today with Yuriy Podolyaka').[13] This graph reveals frequent forwarding relationships between @yurasumy and other prominent pro-Kremlin channels, such as *Resident*, *Oleg Tsaryov*, and *Ostashko! Important*. This level of structured cross-promotion can indicate that channels are working together as part of a coordinated network rather than acting independently.

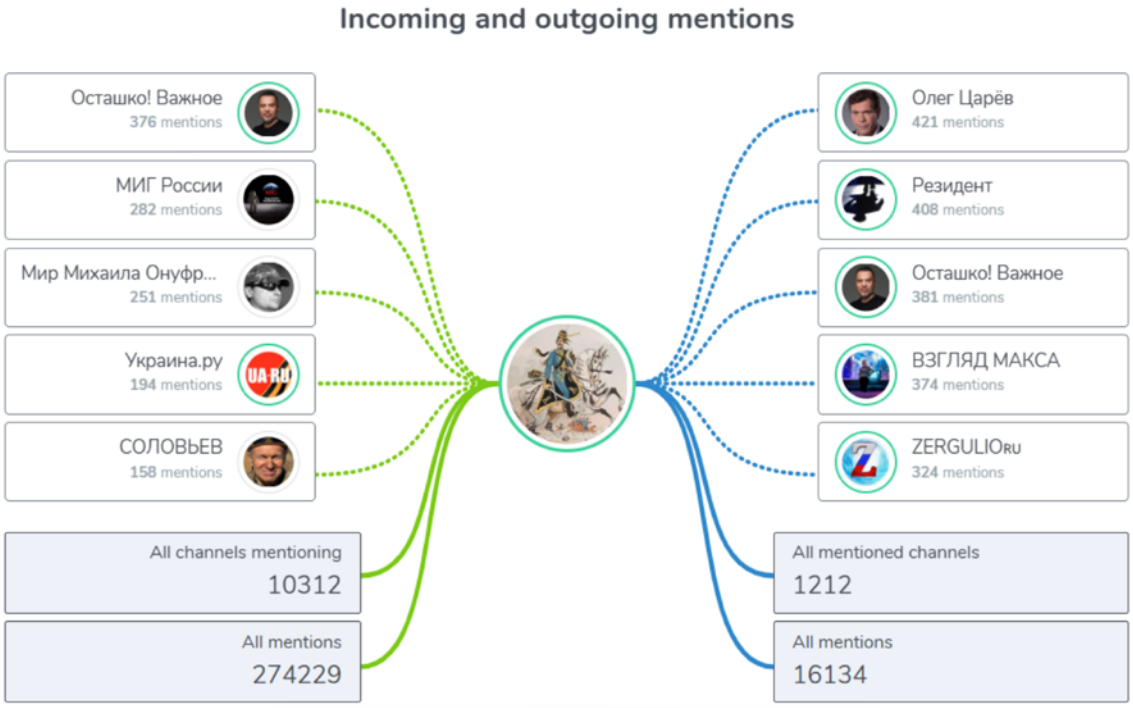

Mention graph for @yurasumy 'The World Today with Yuriy Podolyaka'

Timing metadata helps identify when messages are forwarded and by which accounts, allowing the detection of automated loops or high-frequency relay networks. If multiple channels share a post with identical phrasing and timestamps just minutes apart, this strongly suggests a synchronised dissemination system.

For example, a network of channels branded as regional (e.g. Zaporizhzhia or Dnipro) reposted coordinated narratives tied to Russian-linked channels like *ZLOY Enot* and *lolka UA*, revealing tightly coupled amplification behaviour likely orchestrated across affiliated nodes.[14]

Such technical indicators can establish links between content origin and dissemination architecture. Observing them together may suggest that what appears organic is part of a managed influence system. Alongside behavioural or contextual evidence, they can help build a strong attribution case.

Financial and commercial

This category covers the financial infrastructure that sustains influence operations, including cryptocurrency wallets, advertising identifiers, and payment processors. Such evidence provides some of the clearest and most direct links to operational control and accountability, since funding flows are harder to obscure than content or behavioural patterns.

An assessment typically begins by tracing blockchain transactions to cluster wallets that may be associated with known actors or operational entities. Analysts may use publicly accessible blockchain explorers or commercial analytics tools to monitor flows of funds and identify patterns of reuse.

Parallel to blockchain tracing, researchers examine ad-tech identifiers, such as Google Ad IDs and Meta Business Manager accounts, to identify monetisation channels and ad-purchase behaviours. When legally accessible, payment processor records can also be analysed to reveal account ownership or coordination across multiple services.

Financial evidence therefore often offers the strongest indicators for attribution. However, because many of these signals require proprietary or law enforcement access, it is also the hardest to obtain. In practice, this means that while financial signals can deliver high-confidence findings, they are rarely available in open-source workflows and tend to serve as high-confidence but sparsely available indicators.

Challenges and integration

Despite the advantages of technical evidence, analysts must navigate several obstacles that complicate attribution. Obfuscation and anonymisation techniques such as VPNs, proxy servers, bulletproof hosting, or blockchain-based domains are often employed to mask origin points. Infrastructure may be rapidly created and discarded to avoid long-term detection. Analysts relying solely on open-source data may find themselves constrained compared to those with access to proprietary or classified datasets.

Still, technical indicators provide a systematic and reproducible foundation for identifying infrastructure-level signals. When these indicators are cross-referenced with behavioural patterns (e.g., content dissemination tactics) and contextual evidence (e.g., geopolitical timing or narrative alignment), they strengthen multi-layered attribution and contribute to more credible assessments of influence operations.

Behavioural evidence

Behavioural evidence focuses on patterns of observable activity that reveal how an IIO is managed, deployed, and sometimes concealed. Rather than focusing on what is said, behavioural analysis examines how messages are crafted, amplified, and disseminated.

Behavioural analysis typically begins with the identification of narratives of interest. This often prompts analysts to trace how a narrative emerged and spread. Open-source data from social media platforms and content archives is then collected to build a dataset. Indicators, such as posting time, repetition of phrasing, cross-posting, or shared metadata patterns across accounts, can then be examined.

Cross-posting is a marketing technique involving the placing or syndication of similar content on different websites or social media channels to increase visibility and engagement.[15,16] This tactic is used in influence operations to create the illusion of organic amplification.

Shorter time intervals between postings by suspicious or coordinated accounts may point to Coordinated Inauthentic Behaviour (CIB), making cross-posting an important behavioural signal. For instance, when multiple outlets simultaneously publish a story with identical wording and visuals, or when content appears on aggregator sites before its alleged source, these anomalies may indicate coordinated dissemination or source laundering. Other red flags include sudden growth in engagement, content syndication across multiple outlets without proper attribution, or mirrored branding designed to obscure the origin of the information.

This dataset becomes the foundation for deeper behavioural analysis, interpreting patterns through the lens of **Tactics**, **Techniques**, and **Procedures** (often referred to as **TTPs**). This can help analysts understand how a campaign was organised and what it aimed to achieve.

TTPs refer to the observable behaviours of a threat actor that describe how a campaign is planned and executed in practice.[17]

Tactics are the highest-level description of the behaviour, covering overarching objectives and goals and how they are to be achieved; for example, a tactic may be to attempt to spread disinformation about a battle.

Techniques provide a more detailed description of the behaviour, covering specific activities that support the tactic. For example, a news report may deliberately use realistic video game footage to spread disinformation about a battle to a target audience that receives most of their news through television.

Procedures provide a lower-level, highly detailed description of the behaviour in the context of a technique; e.g., a detailed description of how certain graphical settings, mods, and data capture methods were used to create fake footage and how it was distributed via social sharing features before appearing in mainstream media.

The primary tool for collecting, analysing and cataloguing behavioural evidence is the DISARM Framework.[18] DISARM is a structured framework with approximately 391 specific behaviours, enabling analysts to classify TTPs systematically. It allows for campaign 'fingerprinting' based on repeatable tactics, making it easier to track and compare IIO methods.

The DISARM red framework[19] was developed to support the investigation of threat actors and how they seek to influence the information environment. It has evolved from a sense-making tool that helps to answer the question 'What is happening?', into a way of identifying patterns of observable tactics and techniques, used to support evidence-based assessments.

Ultimately, behavioural analysis provides insight into how IIOs are structured and

executed. It exposes the dynamics behind message spread, including artificial amplification, impersonation, and cross-platform coordination. While behavioural evidence alone rarely confirms attribution, it often highlights the operational logic of influence campaigns, identifying signals that merit deeper technical or contextual follow-up. When used in combination with technical indicators (e.g., infrastructure reuse) or contextual factors (e.g., geopolitical alignment), behavioural evidence can strengthen attribution claims or reveal inconsistencies.

Contextual evidence

Contextual evidence centres on the **content, timing, and geopolitical context** of information influence, helping analysts understand **what** is being said, **how** and **why** it is being said, and **to whom** it is directed.

This evidence focuses on narratives, content delivery, cultural and linguistic indicators (use of dialects, symbols etc.), and alignment with real-world events.

The first step in contextual analysis is identifying a relevant dataset. This typically begins with the detection of suspicious content, often flagged through behavioural patterns (e.g., coordinated timing, copy-paste posting, or bot-like amplification), or keyword monitoring tools that track emerging narratives. Analysts might use tools such as social media listening platforms, Telegram scrapers, or platform APIs to collect relevant posts, hashtags, videos, memes, or URLs tied to the narrative under investigation.

Prioritisation depends on operational goals. Analysts assess reach (engagement levels, virality), coordination indicators (similar content pushed by multiple accounts across platforms), or strategic relevance (alignment with geopolitical events, state media narratives, or known adversarial objectives). For example, during the run-up to Ukraine's 2024 mobilisation reform vote there was an uptick in anti-mobilisation hashtags *(#ТЦК, #stopTRC*)[20] coinciding with viral Telegram content that was then republished on TikTok and Instagram.[21]

This preliminary step establishes the core of the dataset that will be examined for narrative content, audience targeting, and geopolitical relevance.

Narratives are the backbone of contextual analysis. In essence, they are simple stories that shape perceptions and give shortcuts to understanding complex issues. They often express things about identity, community, and purpose. Importantly, they may not be literally true, but rather carry the aggregated, distilled beliefs of a community built up over time by many people across many statements. This includes information about the values, identities and beliefs that drive these narratives, and with whom they have credibility. For instance, Russian propaganda often claims that 'NATO provoked the war in Ukraine', reducing a complex geopolitical conflict into a single blame narrative that resonates with certain anti-Western audiences.[22]

Identifying the narrative and political context involves mapping the content to broader themes: What is being said? Does it reflect a known strategic narrative (such as, anti-NATO, pro-Kremlin, anti-vaccine)? What real-world events, political moments, or social tensions might it be exploiting? Can the narratives be linked to a larger narrative arc, do they speak to locally held grievances or a past or upcoming event?

Understanding the origins of malicious narratives can be done through the concept of **narrative laundering**, a Soviet-era disinformation dissemination strategy, designed to obscure the origins of false narratives.[23,24] Narrative laundering describes how false narratives are legitimised through staged dissemination. The typical stages, **placement** (introducing fabricated material), **layering** (repetition via mixed outlets), and **integration** (amplification through mainstream or foreign channels) were first demonstrated by the KGB's Operation INFEKTION in the 1980s. Modern technologies such as generative AI and cheap digital publishing have intensified these processes, enabling entire networks of fabricated news sites and videos to mass-produce synthetic legitimacy.[25]

Another useful concept is Russia's 'firehose' model of propaganda,[26] where the same ecosystem distributes divergent, even contradictory, versions of events to reach multiple audience segments. For instance, while *Sputnik* in English may frame Ukrainian governance as corrupt and illegitimate, its Spanish-language output has focused on anti-U.S. sentiment and vaccine conspiracies, tailoring narratives to local worldviews. This tactic aims not to persuade with consistency, but to overload and destabilise through volume, variation, and repetition.

Supporting evidence includes identifying media ecosystems and attribution clusters. Analysts distinguish between statecontrolled, statelinked, aligned, and proxy outlets, ranging from overt actors such as *RT* and *Sputnik* to covert content farms funded through oligarch networks and companies like the Social Design Agency.[27] Even if such outlets are not narrative originators, their amplification patterns often serve as contextual indicators of Kremlin proximity.[28]

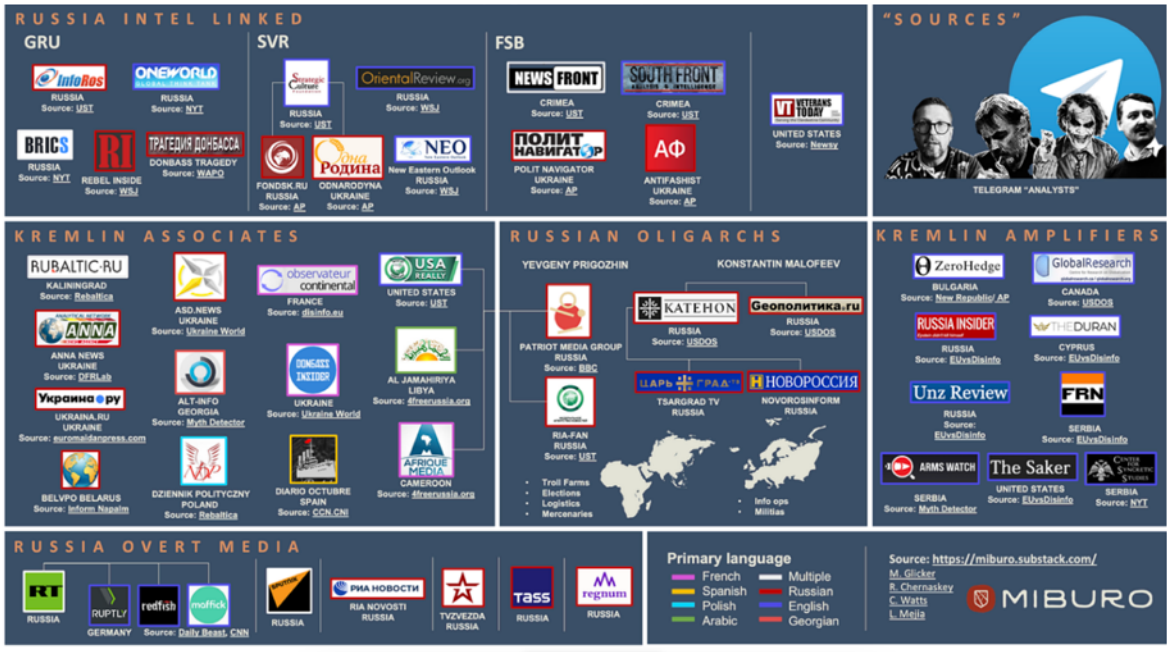

Russia's Information Influence ecosystem (2022)

Analysis of Russia's information influence ecosystem identifies the following categories of influence platforms:

- **Russia intel linked** Media outlets that are linked to Russia's intelligence services, such as GRU, SVR and FSB.
- **Sources** Russia's disinformation ecosystem operates on the content created by online personas such as influencers claiming to be independent experts.
- **Kremlin associates** Outlets that have at one time or another been exposed for having ties with known Russian influence agents, and work to amplify the Kremlin's narratives.
- **Russian oligarchs** Media outlets that have been exposed for having ties to Russian oligarchs, notably the late Yevgeny Prigozhin, and sanctioned businessman Konstantin Malofeev.
- **Kremlin amplifiers** Mostly foreign-facing media outlets, that consistently share news aligned with the Kremlin's objectives, and have been investigated by multiple organisations.
- **Russia overt media** Media outlets that have either open or proven financial ties to the Russian state.

Next, the analyst examines the **form and delivery** of the content. Are the same phrases, hashtags, or slogans repeated across platforms? Are memes or images reused with different captions in different languages?

The investigation then moves to indicators of **audience targeting**. Who benefits and who is the target of a narrative? What languages or dialects are used? Are there regional references, religious or cultural framing, or symbols aimed at a specific demographic? For example, in 2023, Russian propagandists spread narratives alleging that Ukrainian authorities plan to demolish churches belonging to the Ukrainian Orthodox Church – Moscow Patriarchate (UOC-MP).[29] These narratives exploit religious sentiment by framing Kyiv as attacking Orthodox believers, thereby resonating with conservative and religious demographics both inside Ukraine and abroad.

Temporal analysis adds a final layer, whether bursts of content align with key political or diplomatic events. Contextual mapping frequently reveals synchronisation with elections, international summits, or military movements.

In summary, contextual evidence includes, but is not limited to:

- False statements or messages supporting a broader narrative.
- The timing and content of messages or narratives that are related to a significant event, such as an election or major political decision.
- Content that revives historically manipulative tropes or previously deployed disinformation.
- Use of language, dialects, symbols, or hashtags that signal alignment with a particular community.
- Who benefits and who is harmed by the spread of specific narratives.
- The apparent target audience of narratives.
- Similarities between narratives and publicly stated positions of governments and officials.

Building an evidential case using contextual analysis can be challenging. It requires a sound understanding of the technical (e.g., metadata, account origins) and behavioural

(e.g., coordinated activity, amplification patterns) aspects of an influence operation. For example, observing whether a network of accounts posts in synchrony across and between platforms, or maintains unusually regular publishing patterns, strengthens the contextual assessment.

Although technical analysis often provides more definitive evidence and contextual analysis is more interpretive, the latter should not be overlooked. Understanding the narratives and cultural framing that drive influence campaigns is equally important for effective attribution, risk assessment, and response.

Legal and ethical assessment

This category sits outside the analytical process and accompanies evidence collection through continuous ethical and legal reflection. For example, is it ethical for an analyst to join a private social media group under a pseudonym? Acceptability varies by role (journalist, academic or official) and by national legal frameworks.

Attribution can also raise political challenges. For example, an analyst may avoid naming an EU-based individual spreading pro-Kremlin messaging under free expression protections, but direct links to a Russian company could justify disclosure. Governments may withhold public attribution to protect sources, while civil society must weigh the benefits of publicity against risks of reprisals. These judgments are complex and context dependent.

In an environment where relatively few single actors have the resources to conduct an end-to-end investigation that culminates in an attribution, the legal and ethical assessment must also consider the trustworthiness of an inherited attribution and the actor that made it. Social media platforms may make attributions when taking down networks of accounts that violate their policies, but do not provide the technical data that informed their decision. Their scope is also limited to activity on their platforms, meaning that related activity happening on another platform will be out of the scope of their attribution assessment. The lack of transparency also makes it difficult for civil society actors to rate a platform's attribution methodology and evidence.[30]

The field of journalism has developed measures to assess the quality, transparency and trustworthiness of news sources. The US-based internet trust service NewsGuard assesses news and information websites based on nine criteria,[31] while ICFN's Code of Principles guides the work of organizations that "publish non-partisan reports on the accuracy of statements by public figures and prominent institutions and other widely circulated claims related to public interest issues".[32] In 2020, the Atlantic Council's DFR*Lab* created the Foreign Interference Attribution Tracker [33] to rate the credibility, bias, evidence, transparency, and impact of allegations of foreign interference relevant to the 2020 US Presidential election.

Finally, the decision to make an attribution public depends on publicinterest thresholds. In March 2018 in the UK, Sergei and Yulia Skripal were poisoned with Novichok. Two others later fell seriously ill after finding the toxin, disguised as perfume, in a public park bin. Six months later, the UK government released an investigation naming two Russian GRU officers as responsible. Though an extreme case of an attempted assassination of defectors, it shows that governments may disclose detailed public attributions when the threat is serious enough.

Confidence intervals

The complexity of building an attribution means that an assessment is rarely 100% certain. Rather, analysts depend on likelihoods such as the balance of probability.

Confidence intervals are well established in intelligence analysis on the grounds that assessments reflect risk and likelihood rather than absolute facts. For example, analysts may assess hundreds of sources including HUMINT and SIGINT to determine that the likelihood of a terrorist attack in a capital city is currently low. They do not know for certain but rather work on the basis of risk.

Similarly, attributions of influence operations often reflect multiple data points which contribute to an overall picture. That picture is best described through language of probability.

Probability can be expressed in different ways, such as percentages or in words, and there can be misunderstandings between people when interpreting expressions of probability.

The UK Government's probability yardstick combines percentages with descriptions.[34] For example, 25-35% is assessed as 'unlikely', while 80-90% equates to 'highly likely'. The Oasis Open project which governs the use of STIX confidence objects [35] gives users a number of options to assign probability, such as 'low, medium, high', numerical scores and more descriptive language scales such as the Admiralty Credibility Scale,[36] Words of Estimative Probability [37] and the DNI Scale.[38]

Oasis Open provides a 'range of values' between 1 to 100, assigning a score to the different probabilities of each scale, meaning that these scales could be used interchangeably.

The ADAC.io framework and Microsoft DTAC model adopt a harmonised threetier structure of low, medium and high. The key requirement is not which scale is chosen, but that each grade is explicitly defined to prevent

Numeric range	STIX (OASIS Open)	Admiralty Credibility Scale	Words of Estimative Probability (Kent 1964)	DTAC	ADAC.io framework
0–19	Very Low / Implausible	Unreliable / Cannot be judged	Almost certainly not true	Very Low Confidence	Very Low Confidence
20–39	Low / Doubtful	Doubtful / Possibly Unreliable	Unlikely / Improbable	Low Confidence	Low Confidence
40–59	Medium / Possibly True	Fairly Reliable / Possibly True	Even Chance / Roughly Even	Moderate Confidence	Moderate Confidence
60–79	Medium / Probably True	Reliable / Probably True	Likely / Probable	High Confidence	High Confidence
80–100	High / Almost Certain	Completely Reliable / Confirmed	Very Likely / Almost Certain	Very High Confidence	Very High Confidence

Confidence scale comparison table

misinterpretation and ensure consistent analytical language. This reduces human error and ensures consistent interpretation across analysts.

Confidence values within a single dataset may differ. For example, attribution of an official government socialmedia account to a state actor may carry high confidence, while related anonymous amplifiers in the same network may only reach medium or low confidence. Analysts therefore report the confidence range rather than a single score.

Expressing uncertainty transparently, i.e. stating what is known, what remains unverified, and varies, helps prevent overstatement and maintains methodological integrity. Clear articulation of probability ensures that attribution findings are both credible and responsibly communicated.

The Spectrum of State Responsibility

Attributing influence operations to state actors can be further complicated by the range of different relationships that a government can have with those engaging in attempts to influence or interfere abroad. Healey's *Spectrum of State Responsibility* model contends that attribution should focus on the needs of policymakers, who may prioritise identifying blame over determining the attacker.[39] The source of an individual attack is not relevant when considering the most important decisions.

We consider how this spectrum, designed to support the political attribution of cyber-attacks, might be applied to the political attribution of influence operations. We will look at each of the ten stages of the spectrum, considering what evidence might be required to build the case to attribute at each stage.

Significant differences exist between the legal and regulatory frameworks for Information Influence Operations and cybersecurity, as well as in their tactics. Healey argues that nations should be held responsible for cyberattacks originating within their borders, even if they did not actively support or commission the attack. In contrast, fewer international norms exist for assigning blame to states for Information Influence Operations. Unlike cyberattacks, these operations can involve both willing and unwitting participants, creating ambiguity over whether certain narratives reflect free speech, self-interest, or state-driven efforts to influence foreign affairs.

The first two categories in Healey's model consider situations where states have taken action to regulate offensive cyber activity: **state-prohibited** and **state-prohibited-but-inadequate**. State-prohibited refers to states that will act to help stop a third-party attack emanating from its territory or using its infrastructure. According to Healey, while these states cooperate to stop the attack, they still bear some responsibility for the insecure systems that have facilitated the attack. State-prohibited-but-inadequate refers to states with a government that wishes to be cooperative, but for reasons such as a lack of appropriate legislation, or technical skills and tools are unable to do so.

According to the United Nations Trade and Development, 80% of countries worldwide have enacted cybercrime legislation and only 13% of countries have no legislation at all.[40] However, in the context of influence operations, what does state-prohibited mean? Some countries have strict rules limiting freedoms of speech. According to Freedom House's 'Free dom in the world 2025', restrictions on media and freedom of expression intensified for the 19th year in a row. Over two-thirds of countries experienced a deterioration in press conditions, as governments expanded censorship, online surveillance, and the criminalisation of

dissent.[41] In April 2024, Iran sentenced rapper and regime opponent, Toomaj Salehi, to death for 'corruption on earth' although this sentence was overturned in June 2024 and he was subsequently released after serving his prison term. [42] Other countries have passed legislation criminalising disinformation, for example a law in Bangladesh that prohibits spreading 'propaganda' about the country's 1971 war of independence, Belarus has legislation to prosecute those who spread false information online and China has outlawed creating or spreading rumours that 'undermine economic and social order'. [43] Many of these laws have received widespread criticism for being political tools to use against critics and opponents and suppress rights to free speech.

For liberal democracies, who aim to defend the right to freedom of speech, the options to regulate against influence operations have so far been limited to four core areas:

- Illegal speech, for example hate speech, glorification of terrorism and incitement to violence.
- Libel and slander legislation.
- Sanctions applied to foreign-owned media sources, for example the EU ban of RT and Sputnik in 2022.[44]
- Regulations to protect the integrity of electoral processes, such as the EU's Digital Services Act which aims to regulate social media platforms.[45]

In the context of an influence operation, to what extent can a government be held responsible for 'allowing' influence operations to emanate from their borders as a result of their inadequate systems? Is it possible for a state to have adequate systems to prevent influence operations emanating from their borders that do not infringe on civil liberties, such as rights to freedom of expression?

For the following eight stages of the spectrum of state responsibility, it might be helpful to examine what kinds of evidence an analyst might need to make such an attribution in the context of influence operation.

The next stage in the spectrum is **state-ignored**. In the cyber context, this means that the national government is aware, but as a matter of policy does not take official action and may even agree with the objectives of the attacker. It is likely that a significant number of Information Influence Operations are ignored by governments for a variety of reasons: the operation may not violate any national laws, it may be difficult to assess the extent to which foreign and/or domestic actors are involved, it may have gone largely unnoticed and therefore has not reached a threshold for action (and any countermeasure might risk bringing unnecessary attention to the influence operation), or it may indeed align with the objectives of the nation state and is therefore not deemed a threat. As governments have limited resources, and many influence operations go largely unnoticed, some foreign influence operations will have to be tolerated by any nation state.

The following three stages are those in which the state is actively encouraging the influence operation, even if they are not ordering or controlling it. Much of the evidence to make attributions to these stages may well be circumstantial, or require access to proprietary information to prove coordination, particularly in the cases where there are efforts to deliberately conceal connections between the actors and the government.

The first of these is **state-encouraged**, described by Healey as being controlled by third parties but encouraged by national governments as a matter of policy. Behavioural evidence could demonstrate that government officials were endorsing or repeating narratives, while contextual evidence might show alignment with political objectives and that the influence operation reflects state narratives.

Next, **state-shaped**, described as being controlled by third-parties, with the state providing some support, such as informal coordination between government officials and the influence operation. In addition to

the contextual and behavioural evidence mentioned in the previous two stages, this stage could also be evidenced by collecting behavioural evidence of informal coordination with government officials, such as attendance at the same events, or contextual evidence that the narratives of the influence operation change as state narratives evolve. In some cases, it may be possible to identify direct connections between influence operation actors and government officials, either in open source or proprietary information. Direct connections might be evidence of correspondence, or influence operation actors and government officials holding management or board positions at the same organisation.

State-coordinated is defined as a nation state coordinating third-party actors and offering support such as technical or tactical assistance, often covertly. Much of this evidence may only be possible to see with access to proprietary information, such as evidence of information sharing, the existence of procurement contracts, or the provision of technical support. Depending on the nature and scale of the influence operation, behavioural evidence may support an assessment that the actor has access to significant resources, for example if they are able to purchase advertising, or appear to be paying staff to create and distribute content.

The final four stages of Healey's spectrum deal with attributions where the state directly commands and controls the operation, either by using third party proxies or directly employed staff. The first of these is **state-ordered**, where a government uses third-party proxies to conduct the operation on its behalf. In information operations, the evidence required to attribute at this stage is very similar to the previous stage, where the difference is likely to be proprietary evidence proving direct command and control from the nation state, and proprietary or behavioural evidence that the government is supporting the operation with significant technical and/or financial resources.

State-rogue-conducted is different from the other stages in the state-abetting and controlling categories, in that while government officials are directly involved in the commissioning and implementation of the operations, this appears to be happening without the knowledge or approval of the government. This means that the narratives may or may not align with official state narratives and interests, depending on the motivation of the rogue officials. Any coordination with government officials is likely to be informal, because of the unsanctioned nature of this kind of operation. Evidence would be required to demonstrate that such an operation was not officially sanctioned, and this would most likely come from proprietary sources.

The final two stages, **state-executed** and **state-integrated** are similar, both requiring evidence that influence operations are under direct control of the government, using government resources and staff. In both these cases, the government may seek to deliberately conceal its involvement in the influence operation, or the connections may be publicly declared.

Ultimately, applying the spectrum helps policymakers calibrate responses. State-shaped activity may merit diplomatic protest, while state-integrated operations justify sanctions or legal action. Attribution, therefore, is not solely about identifying the operator but about defining an appropriate level of state responsibility and response proportionality within international norms.

State-								
CRITERIA	IGNORED	ENCOURAGED	SHAPED	COORDINATED	ORDERED	ROGUE CONDUCTED	EXECUTED	INTEGRATED
Direct or indirect benefit to state interests	■	■	■	■	■		■	■
IO reflects state narratives		■	■	■	■		■	■
Government officials endorse or repeat narratives		■	■	■	■		■	■
Official state media amplifies IO content			■	■	■		■	■
Informal coordination with government officials			■	■	■	■	■	■
IO narratives change as state narratives evolve			■	■	■		■	■
Informal connections between IO actors and gov't officials			■	■	■		■	■
Direct connection between IO actors and gov't officials				■	■		■	■
Information sharing between IO actor and govt					■		■	■
Official paper trail between IO actors and govt					■		■	■
Government-provided technical support					■		■	■
Direct financial connection					■		■	■
Direct command and control from government					■		■	■
Official government involvement							■	■
Unsanctioned government involvement						■		
Government-provided infrastructure							■	■
Connection between IO actor and government concealed		■	■	■	■			
Overt connection between IO actor and government								■

The chart lists a number of the points to prove each attribution stage. Dark blue = points to prove, light yellow = dependent on circumstances (may not be relevant.)

Attributing Russian influence about Ukraine

This chapter demonstrates how technical, behavioural, and contextual evidence combine in practice to produce attribution assessments.

Each evidence type is examined, with examples to illustrate methods and outputs, then we show how the three evidence types are combined into a final attribution assessment.

TECHNICAL ANALYSIS

Technical analysis provides measurable data points (domains, IPs, timestamps, metadata). However, it's important to recognise that technical evidence on its own acts as the 'sensors', while behavioural and contextual analysis interprets that data.

Readers should not be confused if elements of technical analysis appear behavioural, because while the data itself is technical, its value lies in how it informs broader assessments of coordination, deception, or strategic intent.

This section provides examples of how an analyst can conduct a technical analysis, focusing on two main areas:

- **Digital infrastructure analysis** examines domain records and hosting metadata to identify coordinated networks of websites;
- **Platform and network analysis** uncovers signs of coordination using open source platform data, specifically Telegram.

Analysis of digital infrastructure

The following examples demonstrate how technical analysis of digital infrastructure can be used to support attribution. The first focuses on a single domain, showing how metadata from WHOIS records, hosting information, and SSL logs can surface indicators of coordination. The second illustrates how these methods are applied at scale in a real-world case study, where a network of domains was uncovered as part of a broader effort to circumvent EU sanctions.

Domain analysis of fondfbr.ru

The domain *fondfbr.ru* came to attention during investigations of a viral disinformation narrative on child deportations. It was identified as the origin of a fabricated story alleging that Ukrainian authorities had forcibly deported children to Spain.[46] [47] The domain is publicly associated with the «Фонд борьбы с репрессиями» ('Foundation for Combating Repression'),[48] an organisation established by Yevgeny Prigozhin,[49] late head of the Wagner Group and founder of the Internet Research Agency.

Summary of technical evidence: Fondfbr.ru Case

Evidence Type	Evidence	Details	Analytical Assessment
Technical	Domain Registrant	Registered to "Private Person"	Unusual for NGO; could indicate intent to obscure ownership
Technical	Registrar	REG.RU	Frequently used in Russian IIOs due to lack of oversight
Technical	Hosting Infrastructure	IP: 31.31.196.192; shared with ~1,700 sites	Suggests scalable infrastructure, likely used for coordinated operations
Technical	SSL Metadata	Let's Encrypt, certificates every 90 days	Suggests operational security, anonymity preference
Contextual	Narrative Linkage	Child deportation disinformation	Aligns with pro-Kremlin influence themes
Technical	SSL certificate logs showing regular renewals via Let's Encrypt	Preference for free, identity-anonymising CA	Aligns with anonymity practices in influence operations; suggests deliberate obfuscation
Technical / Contextual	Known Affiliation	Public association with Yevgeny Prigozhin	Ties to Russian state-aligned influence networks

Metadata about the domain, which includes registrar name, registration and expiration dates, name servers, IP address, hosting provider, and ownership information was extracted using the WHOIS tool.[50]

Analysis of this data reveals that fondfbr.ru was registered through REG.RU, a major Russian domain registrar, providing a low-cost, privacy-protected, and anonymous registration service favoured by Russian threat actors. This registrar has reportedly been used in pro-Kremlin media campaigns due to its domestic registration status, which allows it to avoid Western takedown mechanisms.[51]

The data also shows that the registrant is listed as a 'Private Person', which may appear unusual for a domain associated with a non-profit organisation. While not inherently suspicious, this deviates from standard practice, as legitimate NGOs often register domains under their institutional names.[52]

The domain is hosted on the IP address 31.31.196.192, which also hosts 1,700 other websites. Such mass-hosting infrastructure is commonly used in scalable online operations, where multiple domains are launched, managed, or rotated from the same server.[53] While shared IP addresses might be common for small websites,

their use by politically sensitive or state-linked domains can indicate unusual infrastructure patterns.

In addition to WHOIS data, SSL certificate transparency logs were reviewed using the crt.sh database to assess the domain's cryptographic metadata. Certificate Transparency (CT) logs provide a public record of SSL/TLS certificates issued to a domain, offering insights into how actively it has been maintained. The logs for fondfbr.ru show that it has been consistently issued certificates by Let's Encrypt, an automated, free Certificate Authority (CA) that does not conduct identity verification checks, with new certificates appearing approximately every 90 days.[54]

This choice is not unusual but it could indicate a preference for anonymity, as this CA does not perform the identity disclosures required by commercial certificate authorities. In combination with the use of the REG.RU registrar and shared hosting infrastructure, the choice of Let's Encrypt contributes to a broader pattern of leveraging low-cost, minimally regulated services.

Individually, none of these indicators prove coordination or state affiliation. These features can appear in many benign setups. But when observed together, especially in combination with known associations to actors like Prigozhin, they collectively increase confidence that the domain is part of a coordinated Information Influence Operation. The real value of technical analysis here lies in its ability to surface these combinations, prompting deeper investigation. In this case, the technical findings supported a broader attribution assessment that links fondfbr.ru to a Russian-backed influence ecosystem.

Infrastructure tracing – RT and Sputnik

Following the EU's 2022 ban on RT and Sputnik,[55,56] analysts at the Institute for Strategic Dialogue (ISD) launched an investigation to determine whether and how these outlets continued reaching European audiences.[57] It is a practical example of technical analysis which uses open-source tools and infrastructure data.

Summary of Technical Evidence – RT/Sputnik Case

Type	Evidence	Findings	Assessment
Technical	WHOIS data showing coordinated domain registration dates post-EU sanctions	Multiple alternative domains (e.g., actualidad-rt.com) registered shortly after sanctions targeting RT/Sputnik	Indicates deliberate creation of workaround domains to sustain reach despite bans
Technical	DNS records pointing to shared IP addresses and known RT nameservers (e.g., ns1.rttv.ru)	Direct technical link between sanctioned outlets and new domains	Confirms affiliation; unlikely to occur without coordinated planning
Technical	Shared Google Analytics Tracking IDs across multiple domains	Same analytics account managing both known RT domains and circumvention sites	Strong evidence of common operational control
Technical / Contextual	Traffic data showing >85% visits from EU Member States	Workaround sites actively accessed by targeted EU audiences	Demonstrates operational success in sanctions evasion and continued audience penetration
Behavioural / Technical	Observed switching between alternative domains in RT's Spanish-language promotion	Coordinated redirection of audiences to maintain visibility	Shows adaptive tactics consistent with long-term operational planning

The investigation began with a list of known RT and Sputnik domains. Analysts monitored official RT social media accounts to detect the promotion of suspicious new URLs. One such example was *actualidad-rt.com*, a Spanish-language workaround promoted on RT's Twitter accounts. These URLs were flagged as potentially affiliated and passed through infrastructure correlation checks.

To verify whether these domains were technically linked to RT, analysts used WHOIS lookups to extract registration metadata such as creation date to find evidence or coordinated registration of alternative domains. In parallel, DNS records were queried to check whether domains pointed to the same IP addresses or used identical nameservers, including known RT-related entries like *ns1.rttv.ru*. Analysts also searched for shared Google Analytics Tracking IDs (UA codes), which signal when multiple websites report to the same analytics account, an especially useful indicator of common ownership or control.

To ensure accuracy and reduce false positives, each domain had to match at least two technical indicators, such as a shared IP and tracking ID, before being classified as part of RT's extended infrastructure. The resulting network revealed clusters of domains registered in the weeks following the EU sanctions, many using Russian registrars such as RU-CENTER or telecom providers like Megafon. Maltego [58] was used to visualise these infrastructure connections and identify domains that were likely deployed together.

Once technical linkages were established, analysts moved to assess real-world reach. Using SimilarWeb, they retrieved traffic statistics for the flagged domains, including visit volumes, top countries of origin, and referral sources. The data confirmed that many of these workaround websites were not only functional but actively accessed by EU-based audiences. Analysts also noted that many visits came from direct links or social media referrals, particularly from Telegram and Twitter. By tracking changes in promotion behaviour, such as RT's Spanish-language accounts switching between alternative domains, they were able to observe active circumvention strategies over time.

Together, the fondfbr.ru and RT/Sputnik cases demonstrate how layered technical evidence – domain metadata, shared infrastructure, and platform behaviour – forms credible attribution foundations. None of the individual markers is decisive, but their combination under consistent ownership patterns and strategic intent enables medium-to-high attribution within a Russian state aligned information ecosystem.

Platforms and networks analysis

Platform and network analysis can uncover signs of coordination and inauthentic influence using open-source platform data. The first example is an analysis of a pro-Kremlin Telegram channel *@yurasumy* and its amplification patterns and audience engagement. The second looks at an IIO about corruption in Ukraine, analysing repost timing, forwarding structures, and comment activity across multiple channels, to reveal coordinated dissemination. These cases show how technical indicators at both channel and network level help distinguish organic activity from centrally managed influence operations.

Telegram channel @yurasumy – *'Мир сегодня с Юрий Подоляка'* (The World Today with Yuriy Podolyaka)

Operated by a pro-Kremlin commentator Yuriy Podolyaka, this Telegram channel has over 3 million subscribers and an unusually high 55% engagement rate by reach (ERR) score.[59] This data, alongside posting frequency, historical growth, and forwarding graphs, was extracted using the TGStat tool.

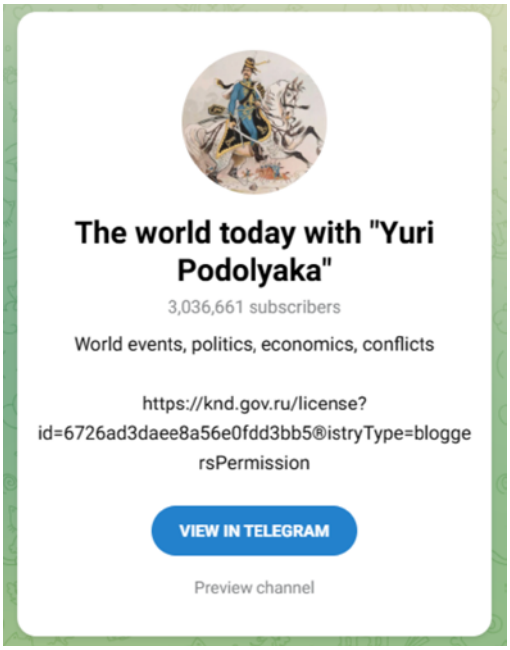

Most striking was the ERR score of 55%.[60] Scores above 50% are generally considered anomalous and may point to artificial view inflation (bots or purchased views), non-organic amplification, or, less commonly, an exceptionally loyal audience.[61]

The channel's amplification network was then mapped using forwarding data and mention graphs provided by TGStat.[62] The data revealed frequent forwarding relationships between @yurasumy and other prominent pro-Kremlin channels, such as *Resident*, *Oleg Tsaryov*, and *Ostashko! Important*. This level of structured cross-promotion often indicates that channels are working together as part of a coordinated network rather than acting independently.

Posting frequency and rhythm were also analysed. The channel produced a high volume of content on a consistent schedule, with few deviations over time. Such regularity may suggest team-based management or semi-automated scheduling.

Finally, the channel was cross-referenced with the official Roskomnadzor (RKN) register,[63] confirming that it was listed among approved Russian media channels. Being listed provides protection under Russian media law and signals formal alignment with state communication objectives

Summary of technical evidence – Telegram channel *@yurasumy*

Evidence Type	Evidence	Finding	Assessment
Technical	ERR score of 55% (above normal range)	Engagement rate significantly exceeds typical organic benchmarks	Suggests possible artificial amplification (bots, purchased views) or highly coordinated promotion
Technical / Behavioural	Frequent forwarding relationships with other pro-Kremlin channels ("Resident," "Oleg Tsaryov," "Ostashko! Important")	Structured cross-promotion within a defined set of channels	Indicates operation as part of a coordinated amplification network
Behavioural	Consistent high-volume posting schedule	Minimal deviation in posting rhythm over time	Implies team-based management or automation in content delivery
Contextual	Inclusion in Roskomnadzor (RKN) register of approved Russian media	Official recognition and protection under Russian media law	Confirms formal alignment with state communication objectives

This shows how technical analysis of platforms and networks can be conducted using the TGStat tool and official listings. Technical indicators like ERR scores, repost timing, and network connections, can be used to determine whether influence is being exerted organically or through centralised, coordinated strategies.

Corruption narrative operation

An IIO about corruption in Ukraine is another example of platform and network analysis that focuses on identifying coordinated content dissemination and message manipulation across several Telegram channels.

A baseline of activity was established using message timestamps,[64] and a subset of 50 incidents revealed that multiple Telegram channels were reposting the same corruption-related content within a short window of 1 to 3 minutes.[65] This narrow repost interval is a strong indicator of automated or pre-scheduled content dissemination.

To understand the network structure behind these messages, network mapping was used to trace where messages originated and how they were forwarded. Many of the posts were traced back to a central source: the Telegram channel *Politika Strany*. These posts were then forwarded by channels mimicking local Ukrainian media outlets. This suggests a deliberate attempt at identity mimicry to increase credibility and local resonance.

The same coordinated content distribution was observed in the comment activity under the posts of high-profile Ukrainian media pages such as TSN and Hromadske. In one incident, 11 accounts published 19 comments using identical or near-identical language to accuse President Zelenskyy and the U.S. of profiting from military aid. These comments matched templates distributed by the Telegram channel *Digital Army of Russia*,[66] a Russian Telegram channel which regularly distributes multilingual comment templates and dissemination instructions.[67,68]

Summary of technical evidence – Corruption narrative operation

Evidence Type	Evidence	Findings	Assessment
Technical	Reposting of identical content within 1–3 minutes across multiple Telegram channels	Highly synchronised dissemination pattern	Strong indicator of automation or pre-scheduled content release
Technical	Network mapping tracing content origin to "Politika Strany" channel	Identifies a single central source behind multiple reposts	Suggests coordinated control over a network of channels
Technical / Behavioural	Forwarding by channels mimicking Ukrainian media outlets	Use of false identities to increase perceived authenticity	Indicates deliberate deception and audience targeting
Technical / Behavioural	Comment flooding with identical/near-identical text on major Ukrainian media pages	19 comments from 11 accounts using templates from "Digital Army of Russia"	Confirms organised use of tasking and template-based messaging

This example demonstrates how repost interval and comment activity analysis and network tracing can be used to indicate automated coordination and centralised messaging strategies.

When examined alongside behavioural evidence, such as the thematic alignment of reposted content, and contextual evidence, such as the political moments chosen for message amplification, these technical findings provide a more complete picture of the operation's intent and structure.

Together, the @yurasumy and corruption narrative analyses demonstrate how platform and network data reveal coordination in proKremlin information ecosystems. High ERR scores, narrowed repost intervals, shared forwarding networks, and identical comment templates provide measurable evidence of centralised control. When these platformlevel indicators are triangulated with behavioural and contextual evidence, such as thematic alignment and timing relative to political events, they enable mediumtohighconfidence attribution linking Telegrambased campaigns to Russian state aligned influence networks.

BEHAVIOURAL ANALYSIS

Behavioural analysis is the examination of patterns and techniques used by IIOs' actors, such as coordination, inauthenticity, and manipulation, to identify how they operate.

This section provides examples of how an analyst can conduct behavioural analysis in support of attribution by examining:

- The use of **cross-posting and source-laundering** techniques in the dissemination of a false story claiming clashes between Georgian and Ukrainian soldiers;
- The **application of the DISARM framework** to classify behaviours and identify patterns behind an influence operation pushing narrative that Poland intends to annex parts of Ukraine.

Cross-posting and source-laundering

Cross-posting has become a widely used method in Information Influence Operations. It involves publishing identical or slightly altered content across multiple platforms to increase visibility, generate the illusion of organic spread, or obscure the original source. To acquire this behavioural evidence, analysts often compare times of postings of the same or similar content across multiple websites and platforms. Shorter posting intervals by fake or suspicious accounts may indicate Coordinated Inauthentic Behaviour (CIB) and reveal the network of malicious actors.

For example, on 25 April, a fake story started circulating on Kremlin-linked media channels about the Georgian and Ukrainian soldiers allegedly attacking each other. It was spread across at least 17 Kremlin-linked outlets, including *Tass.ru*, *Ria Novosti*, *Lenta.ru*, and one Georgian-language source, *Geworld.ge*. Some news outlets cited Tass as a source, while others cited Ria Novosti. Ria, in turn, cited a personal Telegram channel of Andrey Marochko, a known propagandist, as a source. The time of publication across platforms was charted into the timeline of postings shown below.

The times of the postings reveal that Tass.ru published the story before it was published by Marochko himself on his Telegram channel. This sequencing inconsistency suggests central coordination and undermines the authenticity of the alleged source. The three posting times by Lenta.ru, *Rambler.ru* and *Kamchatskoe Vremya* coincide, and share an exact same text and picture. This pattern suggests the use of centralised publishing software to coordinate dissemination. Moreover, several outlets substituted the Telegram source with more 'legitimate' citations (e.g., TASS), demonstrating a source-laundering technique.

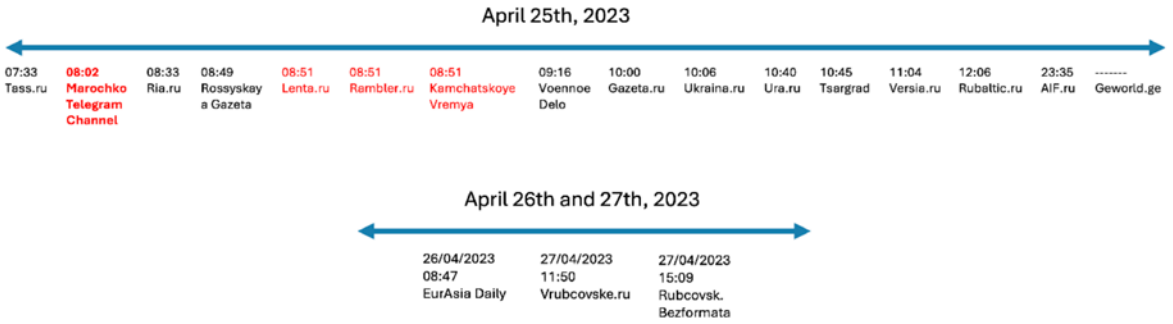

Summary of behavioural evidence – Cross-posting and source laundering

Evidence	Finding	Assessment
Publication of identical or near-identical story across 17 Kremlin-linked outlets	Coordinated, multi-platform dissemination	Strong indicator of a synchronised campaign rather than organic spread
Sequencing anomaly: Tass.ru published before original Telegram "source" (Marochko)	Source chronology inconsistent with claimed origin	Suggests central coordination and undermines credibility of the stated source
Identical text and images across Lenta.ru, Rambler.ru, and Kamchatskoe Vremya	Simultaneous publication with matching content	Implies use of centralised publishing tools
Substitution of original Telegram source with "legitimate" outlets (e.g. TASS)	Source laundering to increase perceived legitimacy	A recognised tactic in Kremlin influence operations

By comparing publication patterns and source attribution across platforms, analysts can build a behavioural profile that reveals a tightly synchronised, multi-channel campaign with obscured origins and amplified messaging, which is consistent with recognised indicators of Coordinated Inauthentic Behaviour (CIB).

While behavioural evidence does not always offer conclusive attribution on its own, it plays a critical role in narrowing the field of actors and highlighting coordinated activity. The following section will explore how such behaviours are classified using the DISARM framework.

Applying the DISARM framework

The following example illustrates how the DISARM framework can be used to identify and document the tactics, techniques, and procedures (TTPs) employed in an Information Influence Operation which appeared during the early stages of Russia's full-scale invasion, when Kremlin-linked Telegram accounts pushed the false narrative that Poland was planning to annex parts of Ukraine.[69] This effort combined forged documents, impersonation of credible sources, and targeted amplification strategies to spread misleading content across platforms and languages.

This case is annotated using DISARM tags to show how specific tactics, techniques, and procedures (TTPs) can be identified and documented. For readability, only key tags are included inline. For a full breakdown of TTP tags used in this case, see Appendix.

Pro-Kremlin Telegram channels continue speculating on alleged Polish aspirations to annex parts of western Ukraine. This narrative was created by circulating fake photos that showed military recruitment posters in the Warsaw Metro station, calling on people to "Stand up for the protection of ancestral Polish lands. Become a Leopard tank operator. Protect Poland in Ukraine".[70] The Russian Telegram channel Signal published forged photos of multiple billboards depicting Jarosław Mika, General Commander of Branches of the Polish Armed Forces, alongside the quote, "It's time to remember history" [71] a reference to the historical fact that parts of western Ukraine were once Polish territory. *(T0086.003: Deceptively Edit Images (Cheap Fakes))*

The channel also mentioned the removal of Ukrainian flags from Polish public transportation, as well as a previous statement by Sergei Naryshkin, the chief of Russia's Foreign Intelligence Service, claiming the US and Poland were plotting to partition Ukraine. *(T0081.004: Identify Existing Fissures)* While the flag removal and Naryshkin's statement both occurred, the channel claimed without evidence that Poland was about to invade Ukraine, using the forged billboards as additional evidence.

The fake billboard story was later amplified by the Kremlin-tied Telegram channel *Gossip Girl*, *(T0098.002: Leverage Existing Inauthentic News Sites)* which had previously published the forged letter alleging Poland's intent to annex Ukrainian territory. "So what does Poland really want?" Gossip Girl asked. "Help secure western Ukraine by sending in troops or regain historical lands?" [72] *(T0102.001: Use Existing Echo Chambers/Filter Bubbles)*

Another Kremlin-tied channel, *Legitimniy* ('Legitimate'), also discussed that Poland might attempt to censor Ukrainian history, ironically reflecting what Kremlin is actually doing. *(T0023.001: Reframe Context)*

Billboard depicting a Polish general alluding to annexing parts of Ukraine

On May 3, Telegram channel *Rokot|Ryk*, which uses the Russian pro-invasion 'Z' symbol in its logo, published a short clip of a speech by Polish President Andrzej Duda in which he said that there wouldn't be any borders between Poland and Ukraine. In its original context, President Duda's quote was in reference to a new era of Ukraine-Poland cooperation and opposition to Russian imperialism and its occupation of Crimea and Eastern Ukraine.[73] Presented out of context, though, the Telegram post misleadingly implied that President Duda was discussing the 'inclusion' of Ukraine within Polish territory *(T0087.002: Deceptively Edit Video (Cheap Fakes), T0023.001: Reframe Context)*.

Pro-Kremlin channels embraced the false interpretation that President Duda intended to annex Ukraine, even suggesting the new country would be renamed 'Ukropol'. On May 5, a video with Russian subtitles appeared on pro-Lukashenka Telegram channel *Zheltye slivi* ('Желтые сливы' or 'Yellow Leaks'). The Signal Telegram channel picked up the video and claimed that in the future, President Duda "would be able to rely on the potential of its neighbours" in the Baltic states to build a community of nations. The channel concluded by describing this as an "imperialist statement," and reiterated that Poland intended to expand its territory. Ukrainian Kremlin-tied channel ZeRada also published the video with the comment, "On what grounds [does Poland] propose to live on Ukrainian land?" and asked whether Ukrainian President Volodymyr Zelenskyy should reply to these claims. The post also alluded to the forged images of Polish billboards.

On May 4, a video displaying the BBC News logo appeared online *(T0087.002: Deceptively Edit Video (Cheap Fakes) (T0097.202: News Outlet Persona, T0143.003: Impersonated Persona))*, *(T0100: Co-Opt Trusted Sources, and T0099: Prepare Assets Impersonating Legitimate Entities)* repeating the same allegation that Poland was preparing to send troops to Western Ukraine "under the pretext of protection from Russia." Captions in the video suggested that the Polish Commander-in-Chief of the Armed Forces had already ordered the army to "prepare for an invasion of Ukraine," which was "confirmed" by a "published order" signed by General Mika. The video also asserted that Washington had endorsed Poland's invasion to Ukraine, while NATO would "officially stand aside".

As evidence, the video included a forged document previously analysed by the DFRLab, which allegedly ordered Polish armed forces to prepare airborne units to be deployed in Ukraine. It also showed the fake billboards of General Mika, as well as a helicopter and Polish soldiers allegedly filmed in northern Poland preparing to deploy to Ukraine. The video was disseminated on Twitter, Telegram, and Facebook in multiple languages including Russian, French, Italian, Turkish and Czech.[74] *(T0101: Create Localised Content)*

Summary of behavioural evidence – DISARM mapping

Evidence	DISARM tags	Findings	Assessment
Forged billboard images of General Mika calling for defence of "ancestral Polish lands"	T0086.003: Deceptively Edit Images	Fabricated visual assets designed to support false territorial claims	Core falsification tactic to lend credibility to fabricated narrative
Exploitation of real events (flag removal, Naryshkin statement) to support false invasion claim	T0081.004: Identify Existing Fissures	Use of factual events to seed plausible but false conclusions	Increases believability by anchoring disinformation in partial truths
Amplification via inauthentic news sites and echo chambers	T0098.002: Leverage Existing Inauthentic News Sites; T0102.001: Use Existing Echo Chambers/Filter Bubbles	Dissemination through Kremlin-linked channels and closed networks	Demonstrates coordinated amplification across known influence assets
Misrepresentation of President Duda's statement via deceptive video editing	T0087.002: Deceptively Edit Video; T0023.001: Reframe Context	Original quote reframed to imply territorial ambitions	Distorts meaning to fit Kremlin narrative objectives
BBC-branded fake video with forged document alleging Polish invasion orders	T0097.202: News Outlet Persona; T0143.003: Impersonated Persona; T0100: Co-Opt Trusted Sources; T0099: Prepare Assets Impersonating Legitimate Entities	Fabrication of authoritative-looking content to legitimise false claims	High-complexity impersonation tactic to mislead audiences
Multi-language dissemination (Russian, French, Italian, Turkish, Czech)	T0101: Create Localised Content	Tailored messaging for multiple linguistic audiences	Expands reach and resonance across target groups

This case demonstrates how DISARM can be applied to map complex influence operations by categorising distinct behaviours, such as content falsification, impersonation, and multi-platform amplification. By documenting these TTPs consistently, analysts can expose coordinated inauthentic behaviour, trace operational fingerprints across campaigns, and improve attribution. Even when the content is misleading or fabricated, the behavioural patterns themselves often reveal intent, coordination, and recurring techniques across actors and narratives.

CONTEXTUAL ANALYSIS

This section explores how contextual analysis can support attribution by looking beyond surface-level content to assess narratives, timing, delivery, audience targeting, and strategic alignment. First, we focus on using **narratives** as a departure point for further research. Second, we apply the **narrative laundering** concept to identify the source of the story about Olena Zelenska allegedly spending a million dollars at a luxury brand store. Finally, we highlight cases where insights emerge **not from the narrative** itself, but from **how and when** it is deployed, offering a fuller picture of operational intent and attribution.

Data collection using narratives

Many approaches to attribution begin by identifying narratives that suggest Kremlin involvement. For example, in their work on influence operations exploiting the issue of corruption, the CSC-IS, together with Osavul, analysed over 130,000 corruption-related messages in Ukraine's online space, collected between July and December 2023.[75] The dataset covered Twitter/X, Telegram, Facebook, YouTube, and online media, with posts in both Ukrainian and Russian. From this baseline, analysts identified 86 information incidents,[76] each consisting of at least six messages clustered around a shared theme. Among the sources active in spreading corruption-related content, 418 had previously been linked to information operations, 462 were directly affiliated with Russia, and 223 were identified as bot accounts.

Information incidents were identified as Pro-Kremlin based on their narratives and the channels that shared them. Pro-Kremlin *channels* were defined as those that systematically disseminate narratives and messages consonant with the central line of Kremlin propaganda and disinformation (Russian state media and known proxies). Pro-Kremlin *narratives* were defined as those which are consonant with the central line of Kremlin propaganda and disinformation.

Seven recurring pro-Kremlin narratives were extracted:

- Zelenskyy is covering up corruption / Zelenskyy himself is corrupt.
- Ukraine is a completely corrupt country.
- Corruption will cause Ukraine to lose the war and alienate the West.
- Ukraine is tied to Western elites through corruption schemes.
- Calls to overthrow the government due to corruption.
- Elites profit while ordinary soldiers suffer.
- Ukraine resells Western weapons.

Each of these narratives was amplified through manipulation tactics such as fake stories, conspiracies, twisting of legitimate reports, and the coordinated distribution of identical posts across Telegram and Facebook.

The data was processed through the Osavul AI-driven platform and supplemented by manual investigation. Analysts cross-checked suspicious content against

open-source evidence and official Ukrainian and foreign sources, ensuring that pro-Kremlin alignment was identified through both narrative framing and source behaviour.

The resulting dataset provides leads for further technical (e.g., bot detection, network mapping) and behavioural (e.g., coordination patterns) analysis, which can ultimately support attribution.

Narrative laundering

Reliable attribution depends on successfully identifying the origins and sources of malicious narratives, which are often intentionally concealed. As outlined in Chapter 2, the concept of 'narrative laundering' involves three key steps: **Placement, Layering and Integration.**

One recent narrative laundering case concerns Ukraine's First Lady, Olena Zelenska, and an alleged million-dollar purchase from a luxury brand store. The story pushed the idea that Ukraine was squandering its Western aid, portraying it as a corrupt nation, not quite 'Brave Ukraine'.[77] It likely resonated with audiences already inclined to think negatively about Ukraine. Within several days, it gathered thousands of reposts and millions of views across multiple platforms.

Summary of contextual evidence – Narrative laundering

Evidence	Finding	Assessment
Three-phase laundering process: Placement → Layering → Integration	Fabricated story seeded in inauthentic accounts, amplified across fake/credible sources, integrated into Western-facing media	Classic narrative laundering model designed to obscure origin and increase credibility
Placement Video of woman claiming Cartier purchase by Olena Zelenska posted to private Instagram	Fabricated witness and forged receipt as initial 'evidence'	Initial seeding point, exploiting perceived insider credibility
Layering Repost by inactive YouTube account, spread across Telegram channels mimicking Russian and Ukrainian sources	Mixed-source amplification to blur origin	Ensures narrative visibility and perceived legitimacy
Integration Story picked up by Russian-aligned media as proof of Ukrainian corruption	Presented to audiences predisposed to anti-Ukraine sentiment	Designed to reinforce pre-existing bias and broader corruption narrative
Investigation revealed woman was a student from St. Petersburg with no ties to Cartier	Debunked core claim	Confirms falsification and deliberate deception

At the end of September 2023, a video was planted in a private Instagram account, allegedly a former Cartier intern, who helped Olena Zelenska with a $1.1 million purchase and then subsequently fired. The woman showed a receipt of the purchase as proof. The video was reposted by a YouTube account with no prior activity, quickly spreading across websites and Telegram channels that mimicked both Russian and Ukrainian sources. An investigation by the Italian news site *Open* later revealed that the woman in the video was a student from Saint Petersburg, Russia, with no ties to Cartier.[78] Despite this, the story was repeatedly cited across Russian-aligned channels as proof of Ukraine's elites being corrupt.

This example shows why narrative laundering is a key concern for attributing narratives. As laundering adapts (through AI, fake accounts, and cross-platform tactics) contextual analysis can help uncover intent and origin.

Looking beyond narratives

Contextual analysis gains power when it expands beyond **what** is said to ask **when**, **how**, and **to whom** it is said. Using two examples below, we will demonstrate how timing, audience targeting, and content delivery can reveal strategic intent and attribution.

Summary of contextual evidence – 'Corrupt Ukraine' narratives

Evidence	Finding	Assessment
Spikes in corruption-related messaging (86 incidents) after key political events (e.g., Zelenskyy's US visit)	Volume and timing aligned with moments of heightened scrutiny	Shows deliberate exploitation of geopolitical context and audience predisposition
Recycling of long-standing Kremlin corruption tropes	Repurposing of familiar narratives for coherence and credibility	Fits established Russian IIO patterns identified in prior campaigns
TikTok anti-mobilisation campaigns timed to Zelenskyy's constitutional mandate expiration and viral civilian–officer encounters	Content timed to exploit domestic political tensions	Indicates opportunistic targeting of sensitive issues
Platform-specific tactics: AI-generated clips, staged encounters, influencer monologues, humour	Content tailored to TikTok's user base (18–55) for emotional impact	Demonstrates audience-specific operational design
Cross-platform monitoring reveals thematic and temporal synchronisation	Patterns not visible in narrative analysis alone	Reinforces attribution by linking timing, format, and targeting strategy

From July to December 2023, baseline monitoring of Ukraine's digital space identified over 130,000 corruption-related posts across X, Telegram, Facebook, YouTube, and local media.[79] Clusters of messages were then identified (minimum six per incident) that pushed coherent corruption narratives. 86 such incidents were flagged as probable components of coordinated influence operations.[80]

This topic was already recognised as particularly sensitive in Ukrainian society, for whom corruption was ranked as a top concern,[81] with international scrutiny intensifying, especially regarding transparency of Western aid.[82] So, the claim targeted audiences already predisposed to view Ukrainian elites with suspicion.

The influence operation mapped onto this environment with precision. The volume and frequency of corruption-related messaging spiked following President Zelenskyy's visit to the United States in December 2023. False claims surfaced alleging American senators had demanded transparency and that Zelenskyy failed to comply. This case repurposes a long-running Kremlin trope that Ukrainian elites misuse Western aid. It recycles historically manipulative narratives to build coherence and familiarity, exploiting vulnerabilities in political trust and attempting to fuel fatigue in Western donor countries. It highlights the importance of temporal and geopolitical context in contextual analysis.

In a different operation that exploited TikTok, analysts also looked beyond the narratives that were being spread.[83] Spikes in anti-mobilisation hashtags (*#ТЦК*, *#stopTRC)*[84] and claims about President Zelenskyy's legitimacy were tightly linked to specific political events such as the expiration of his constitutional mandate in May 2024 and encounters between civilians and mobilisation officers which went viral.

The use of TikTok is also significant, as the platform increasingly overtakes other forms of media among Ukrainians aged 18–55.[85] Rather than presenting detailed arguments, the campaign relied on the form and delivery of content to provoke emotional reactions: AI-generated news clips, staged confrontations, influencer-style monologues, and platform-specific humour, techniques well-matched to the platform and its audience.

In both cases, contextual analysis provided insights into the timing, tailoring, and impact of the operation that would not have been apparent through narrative analysis alone. Recognising these broader indicators allows analysts to attribute campaigns more accurately and to detect emergent threats before they achieve scale.

Integration and final assessment

Here, we conduct an attribution assessment of the operation related to the corruption narrative campaign already examined in earlier sections,[86] expanding the analysis by bringing technical, behavioural, and contextual evidence together, and applying the state responsibility framework and confidence levels to create an integrated attribution assessment.

	Open source evidence	Proprietary evidence	Confidence level
Technical	Linked Telegram channel ownership; botnets; simultaneous posting	Osavul platform identified 462 Russian-affiliated sources and 223 bots repeatedly spreading corruption narratives	Moderate–High
Behavioural	Coordinated posting, tasking by 'Digital Army of Russia', use of fake-Ukrainian personas, tactic shifting	N/A	High
Contextual	Narratives mirror Russian state messaging, timed with geopolitical events, amplified by state outlets	N/A	High
Gaps	No direct proprietary data (e.g. payments/contracts), some amplification by unwitting/non-state actors		
State responsibility	State-shaped to state-coordinated		High confidence
Overall Confidence	Strong convergence across evidence classes		High (≥80%)
Attribution	It is assessed with high confidence that this IIO is state-shaped to state-coordinated by the Russian Federation.		

Bringing together the technical, behavioural, and contextual evidence helps to pinpoint the operation's orchestrators instead of merely listing separate indicators.

Technical analysis reveals interconnected networks of Telegram channels and bot accounts that routinely posted coordinated content within minutes of each other. Monitoring flagged hundreds of Russian-affiliated sources and bots distributing similar corruption narratives. These digital traces demonstrate a pattern of organised and repeated activity, suggesting a managed and sustained campaign rather than isolated or spontaneous acts.

Behavioural evidence builds on this by demonstrating the intent and operational patterns exhibited by the actors involved, such as using impersonation of Ukrainian media and coordinated comment flooding campaigns to embed the messages in domestic discourse. Specific elements such as 'tasking' instructions from known Russian coordination hubs like the Digital Army of Russia provide clear signals of centralised coordination and purposeful amplification. Near-simultaneous publishing across channels connected to known pro-Kremlin clusters, including those tied to channels *ZLOY Enot* and *Iolka UA*, further indicates direction and control consistent with state involvement.

Contextual assessment situates these patterns within the broader political landscape, underscoring alignment with Russian state objectives. As outlined in the previous section, narratives align closely with Russian state propaganda, aiming to undermine trust in Ukrainian

institutions, weaken international support, and exploit societal divisions. Moreover, bursts of activity correlate with moments critical to Russia's geopolitical interests. The use of government-aligned media channels to propagate these messages further reinforces the context of state sponsorship and strategic intent.

By combining these three evidence types, analysts move beyond individual data points to a whole understanding of the operation. Technical data reveals the 'how' of the campaign's execution; behavioural patterns define the 'what' of its actors and methods; and contextual analysis explains the 'why' grounded in political objectives. While the open-source nature of this investigation limits access to proprietary data such as internal tasking orders or financial records, the available evidence is sufficient to rule out purely organic or rogue-actor explanations.

Regarding responsibility, the operation's characteristics place it between state-shaped and state-coordinated. There is a potential third-party control through networks and bots, but with substantial informal to formal support by state actors. It is evidenced in state-aligned media amplification, adaptive narrative shifts mirroring official rhetoric, and resource allocation such as infrastructure and botnet usage.

While open-source material does not prove direct command or ownership at the levels of state-ordered to state-executed, there is substantial behavioural and contextual alignment with Russian state interests. There is no indication that the Russian government is unaware of these activities, nor any sign of prohibition or corrective action. The patterns are also inconsistent with operations run by rogue officials, as the activity is overtly aligned with Russian national objectives and promoted through official channels.

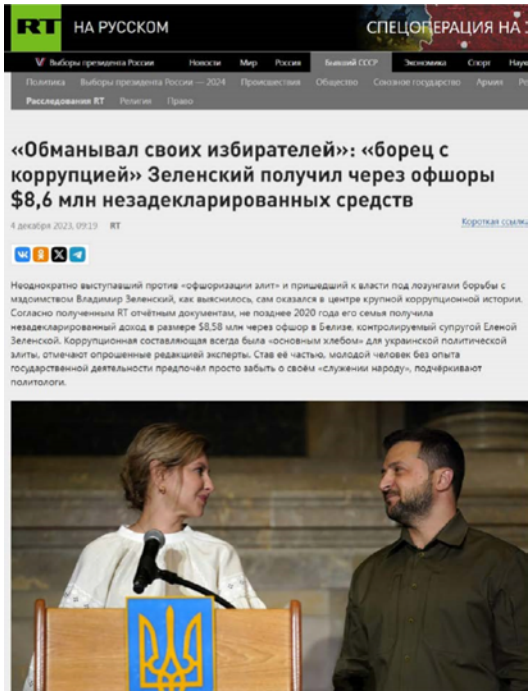
RT НА РУССКОМ СПЕЦОПЕРАЦИЯ НА

«Обманывал своих избирателей»: «борец с коррупцией» Зеленский получил через офшоры $8,6 млн незадекларированных средств

4 декабря 2023, 09:19 RT

Неоднократно выступавший против «офшоризации элит» и пришедший к власти под лозунгами борьбы с мздоимством Владимир Зеленский, как выяснилось, сам оказался в центре крупной коррупционной истории. Согласно полученным RT отчётным документам, не позднее 2020 года его семья получила незадекларированный доход в размере $8,58 млн через офшор в Белизе, контролируемый супругой Еленой Зеленской. Коррупционная составляющая всегда была «основным хлебом» для украинской политической элиты, отмечают опрошенные редакцией эксперты. Став её частью, молодой человек без опыта государственной деятельности предпочёл просто забыть о своём «служении народу», подчёркивают политологи.

RT article alleging Zelensky had smuggled weapons

Confidence in this attribution is rated high (>80% likelihood) due to the strong convergence of multiple, independent evidence classes. While a gap remains in proprietary technical data like payment records or internal state documents, this absence does not undermine the high likelihood assessment given the overwhelming joint behavioural and contextual corroboration supported by technical signals.

Conclusions and Recommendations

This report has examined how Russian Information Influence Operations (IIOs) can be attributed using an Information Influence Attribution Framework (IIAF) that combines technical, behavioural, and contextual evidence. Through case studies we have tested how the IIAF performs in practice under predominantly opensource conditions and where it needs refinement. Three main conclusions arise.

First, attribution of IIOs differs from cyber attribution. While cyber forensics often rely on controlled access to logs, malware signatures, and classified intelligence, IIO attribution is built largely on open-source data of variable quality. In our domain and platform cases, technical indicators such as WHOIS records and SSL certificates provided leads, but they were fragmentary, requiring greater weight on behavioural and contextual analysis to build a credible picture.

Second, no single line of evidence is sufficient. Across the cases, the IIAF helped clarify practical evidential thresholds for high-confidence attribution in largely opensource environments. In practice, robust attributions required independent indicators from at least two evidence categories, explicit documentation of residual gaps, and stated confidence ranges, rather than categorical claims of certainty.

The corruption narrative case provides the clearest demonstration of how technical, behavioural, and contextual evidence must converge to support high confidence attribution. Technical indicators, such as channel metadata and posting anomalies, suggested inauthentic activity but were inconclusive in isolation. Behavioural analysis of synchronised reposting across a Telegram cluster revealed coordination inconsistent with organic dissemination, while contextual analysis situated these behaviours within a longrunning Kremlin corruption narrative exploiting Ukrainian concerns about elite misconduct. Taken together, these strands enabled a structured assessment that the campaign was proKremlin and fell between stateshaped and statedirected involvement, illustrating the value of weighing multiple categories of evidence in parallel rather than seeking a single ‘smoking gun’.

A third conclusion is the value of standardisation of language and frameworks. The application of the DISARM framework demonstrates how standardised categorisation of tactics and techniques can expose operational fingerprints. Greater adoption of such frameworks would support more consistent attribution across governments, platforms, and civil society.

Attribution decisions have both technical and political dimensions. Technically, attribution may risk exposing covert sources or enabling adversaries to adapt. Politically, governments weigh attributions against geopolitical strategy, domestic politics, and available political capital. In some cases, political leaders may avoid attribution despite strong evidence; in others, they may pursue it even when the evidence base is weaker. This dynamic interplay is not always rational but must be recognised as central to the practice of attribution.

Improving attribution practice

Transparency and standardisation Attribution should document confidence levels and limitations explicitly, using clear probabilistic language. Chapter 3 demonstrated how structured summaries can communicate both evidence and confidence concisely. To operationalise this, organisations conducting public attributions should adopt a shared confidence scale (for example, adapted from the UK Probability Yardstick or STIX confidence objects) and require that every major claim is accompanied by: (1) a stated confidence range, (2) a brief explanation of which evidence categories it rests on, and (3) a short rationale for key gaps or caveats. This type of templated reporting would make open-source assessments more comparable across governments, platforms, and civilsociety actors and easier to defend in regulatory or judicial settings

Access to proprietary and classified data Current reliance on partial open-source and platform disclosures limits robustness. Secure, vetted mechanisms for sharing sensitive data could strengthen public attribution. Enhanced cooperation between governments, platforms, civil society, and independent researchers is essential to close gaps. In practice, this implies developing tiered access models, for example, trusted research environments or data-safe havens, where vetted investigators can query platform telemetry, ad-tech records, or law-enforcement data without unrestricted bulk export. Clear protocols for how such restricted evidence feeds into publicfacing IIAF assessments (e.g. noting that a conclusion is supported by classified or proprietary sources without disclosing details) would help reconcile transparency with source protection.

Refined attribution language Greater precision in describing degrees of state involvement (e.g., 'state-shaped,' 'state-integrated') would reduce ambiguity and improve accountability. Analysts should anchor their wording in the spectrum of state responsibility used in this report, explicitly stating which stage an operation is assessed to fall into and why. Over time, adopting a small, shared glossary for terms such as state-encouraged, state-shaped, state-coordinated, and state-integrated would support more consistent sanctions decisions, contentmoderation actions, and strategic communications across different institutions.

Anticipatory analysis Most attribution is retrospective. Developing predictive and real-time analytical capabilities would help detect emerging campaigns before they achieve impact. Building on the corruptionnarrative monitoring work, one practical step is to institutionalise continuous baseline tracking of sensitive themes (e.g. corruption, mobilisation, territorial integrity) and to flag anomalous spikes, crossplatform synchronisation, or rapid narrative laundering as earlywarning signals. Integrating such alerts with the IIAF, so that suspicious activity is quickly triaged for technical, behavioural, and contextual analysis, would shorten the time from detection to defensible attribution.

Systematised network and TTP analysis Automated tools for mapping reposting networks, bot activity, and operational tactics are increasingly necessary to match the scale of adversary operations. The report's Telegram and DISARM case studies suggest that this systematisation should include: (1) routine generation of repost and mention graphs for priority channels, (2) automated detection of anomalous engagement metrics such as extreme ERR scores, and (3) consistent tagging of tactics, techniques, and procedures in line with DISARM or similar taxonomies. Embedding these processes into opensource workflows would make it easier to reuse prior knowledge about actors and campaigns when new operations emerge.

Documented methodological workflows Analysts should publish high-level descriptions of their investigative workflows, including narrative selection criteria,

infrastructure tracing steps, DISARM coding decisions, and integration procedures, so that IIAF-based attributions can be understood, critiqued, and replicated by other actors. This kind of procedural transparency is particularly important when evidential thresholds may be tested in regulatory or judicial settings.

Capacity building for non-state investigators Targeted training, tooling support, and legal guidance should be developed for civil society organisations, journalists, and research institutes, which provide much of the open-source evidence in Ukraine-related cases but often lack access to standardised methods and secure data-sharing mechanisms. Strengthening their capacity is essential to maintaining a diverse, resilient attribution ecosystem that does not rely solely on governments or platforms.

Attribution of IIOs remains challenging, but the IIAF demonstrated in this report shows that highconfidence assessments are possible when technical, behavioural, and contextual evidence converge and are expressed through transparent confidence and stateresponsibility scales.

Future progress will depend on increased transparency, greater standardisation, and closer collaboration between stakeholders. Attribution will remain as much a political act as an analytical one. However, by strengthening the evidentiary base, refining our language, and improving our methods, we can ensure that attribution remains credible, actionable, and resilient in the face of adversary adaptation.

Endnotes

1 A Russian Information Influence operation is a planned, coordinated campaign by Russian state or state-aligned entities to manipulate audiences through deceptive or coercive information tactics, aiming to advance Russian strategic objectives by undermining truth, trust, and democratic coherence in target societies. See Palmertz, B., Isaksson, E., & Pamment, J. (2025). A framework for attribution of information influence operations (ADAC. io Deliverable D1.1). Psychological Defence Research Institute, Lund University.

2 The strategic use (or threat) of legal action and litigation as a tool of information influence, aimed at pressuring, deterring, or punishing regulators, platforms, journalists, or researchers who expose or restrict Russian statelinked operations

3 The Guardian. (2024, June 2). **Revealed: Russian legal foundation linked to Kremlin activities in Europe.** The Guardian. Retrieved May 7, 2025.

4 Iyer, P. (2024, August 4). **Researchers consider the impact of Meta's CrowdTangle shutdown.** *Tech Policy Press*. Retrieved May 7, 2025; Robertson, A. (2023, May 31). **Twitter just closed the book on academic research.** *The Verge*. Retrieved May 7, 2025.

5 See Pamment, J., & Smith, V. (2022, July 19). *Attributing information influence operations: Identifying those responsible for malicious behaviour online.* **NATO Strategic Communications Centre of Excellence & European Centre of Excellence for Countering Hybrid Threats.** Retrieved May 7, 2025, and Palmertz, B., Isaksson, E., & Pamment, J. (2025). ADAC.io Attribution Framework Report (Deliverable D1.1): **Attribution methodology for information influence operations, evidence assessment, and case studies.** ADAC.io Project, Lund University Psychological Defence Research Institute.

6 European Parliament & Council of the European Union. (2022, October 27). **Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market for Digital Services and amending Directive 2000/31/EC (Digital Services Act).** *Official Journal of the European Union*, L 277, 1–102. Retrieved May 7, 2025.

7 Carnegie Endowment for International Peace. (2021, January 26). **What is so foreign about foreign influence operations?** Carnegie Endowment for International Peace. Retrieved May 7, 2025.

8 Pamment, J., & Smith, V. (2022, July 19). **Attributing information influence operations: Identifying those responsible for malicious behaviour online.** NATO Strategic Communications Centre of Excellence & European Centre of Excellence for Countering Hybrid Threats. Retrieved May 7, 2025.

9 Palmertz, B., Isaksson, E., & Pamment, J. (2025). **ADAC.io Attribution Framework Report (Deliverable D1.1): Attribution methodology for information influence operations, evidence assessment, and case studies.** ADAC.io Project, Lund University Psychological Defence Research Institute.

10 Ibid, p15.

11 **WHOIS Record for FondfBr.ru**

12 **The TGStat page for 'The World Today with Yuriy Podolyaka'.**

13 **Forwarding data and mention graphs provided by TGStat for "The World Today with Yuriy Podolyaka".**

14 Center for Strategic Communications and Information Security. (2024, April 16). **"How Russian propaganda speculates on the topic of corruption in Ukraine".** (Retrieved June 16, 2025)

15 Later. (n.d.). **What is cross-posting? In Social media glossary.** Retrieved October 23, 2025.

16 IndieWeb. (n.d.). **Cross-posting.** Retrieved October 23, 2025.

17 National Institute of Standards and Technology. (n.d.). **tactics, techniques, and procedures (TTP) – Glossary | CSRC.** Retrieved May 7, 2025.

18 DISARM Foundation. (n.d.). **DISARM Red Framework.** Retrieved May 7, 2025.

19 The DISARM Red framework helps analysts identify and describe the TTPs used by actors conducting influence operations; it complements the Blue framework, which focuses on defensive responses.

20 ТЦК: The Territorial Center of Recruitment and Social Support (TRC) is Ukraine's military administration body that keeps military records and mobilises the population.

21 ТЦК: The Territorial Center of Recruitment and Social Support (TRC) is Ukraine's military administration body that keeps military records and mobilises the population.

22 Centre for Strategic Communications and Information Security. (2024, November 20). **"Exploiting TikTok for malicious influence on ukrainian audience".** (Retrieved July 17, 2025)

23 Centre for Strategic Communications and Information Security. (2023, May 25). **"Narratives of Russian Propaganda Common in Neighbouring Countries of Ukraine".** (Retrieved July 21, 2025)

24 Linvill, D., & Warren, P. (2023). **Infektion's evolution: Digital technologies and narrative laundering. Media Forensics Hub Reports,** (Report No. 3). Clemson University.

25 See Korta, S. (2018). **Fake news, conspiracy theories, and lies: An information laundering model for homeland security (Master's thesis).** Naval Postgraduate School.

26 Meleshevich, K., & Schafer, B. (2018). **Online information laundering: The role of social media.** Alliance for Securing Democracy, German Marshall Fund of the United States.

27 Paul, C., & Matthews, M. (2016). **The Russian "Firehose of Falsehood" Propaganda Model: Why It Might Work and Options to Counter It.** RAND Corporation.

28 European External Action Service. (2025, March). **3rd report on foreign information manipulation and interference threats: Exposing the architecture of FIMI operations** (pp. 42–44). Strategic Communication and Foresight, SG.STRAT, EEAS.

29 Detector Media. (2024, February 5). **Fake: In Kyiv, more than 70 canonical churches of the UOC can allegedly be destroyed.** (Retrieved July 21, 2025)

30 Pamment, J., & Ahonen, A. (2023). **The ethics of outsourcing information conflict: Outlining the responsibilities of government funders to their civil society partners.** NATO Strategic Communications Centre of Excellence. Retrieved May 7, 2025.

31 NewsGuard. (2025, January 22). **Website rating process and criteria.** Retrieved May 7, 2025.

32 International Fact-Checking Network. (n.d.). **The commitments of the Code of Principles.** Retrieved May 7, 2025.

33 Foreign Interference Attribution Tracker, Accessed 23 July 2024.

34 Defence Intelligence – communicating probability, UK Government, 17 January 2023.

35 Oasis Open Standard, STIX Version 2.1, 10 June 2021, Appendix A.

36 MISP Project. (n.d.). MISP taxonomies: Admiralty scale. Retrieved May 7, 2025

37 Kent, S. (1964). Words of estimative probability. Studies in Intelligence, 8(4). Central Intelligence Agency. Retrieved May 7, 2025.

38 See Irwin, D., & Mandel, D. R. (2018). How intelligence organizations communicate confidence (unclearly) (DRDC-RDDC-2021-N121, NATO SAS114 Final Report, Chapter 19). Defence Research and Development Canada. Retrieved May 7, 2025.

39 Healey, J. (2012). Beyond attribution: Seeking national responsibility for cyber attacks. Atlantic Council. Retrieved May 7, 2025.

40 United Nations Conference on Trade and Development. (n.d.). Cybercrime legislation worldwide. Retrieved May 7, 2025; Statista. (2024, June). Share of countries worldwide having active cybercrime legislation in place as of June 2024. Retrieved May 7, 2025.

41 Freedom House. (2025). Freedom in the world 2025: The uphill battle to safeguard rights. Freedom House.

42 BBC News. (2024, June 22). Iranian rapper's death sentence overturned – Toomaj Saleh.

43 Human Rights Watch. (2013, September 13). China: Draconian legal interpretation threatens online freedom. Retrieved May 7, 2025.

44 Radio Free Europe/Radio Liberty. (2024, February 3). Two years into EU ban, Russia's RT and Sputnik are still accessible across the EU. Retrieved May 7, 2025.

45 EU Digital Services Act. (2024, November 4). Digital Services Act (DSA) | Updates, compliance. Retrieved May 7, 2025.

46 Center for Strategic Communications and Information Security. (2025, February 27). "Ukrainian Crimes Against Children": The Structure of the Russian Narratives and Information Operations. (Retrieved June 16, 2025)

47 Foundation for Combating Repression. FOUNDATION FOR COMBATING REPRESSION FOUND EXCLUSIVE EVIDENCE OF TRAFFICKING DISABLED UKRAINIAN CHILDREN IN SPAIN (Retrieved June 30, 2025)

48 The organisation's official English name, as stated on its website, is 'Foundation to Battle Injustice'. However, for this report it is referred to as the 'Foundation for Combating Repression', a translation that more accurately reflects the original Russian name.

49 Vedomosti. (2021, March 23). Пригожин основал «Фонд борьбы с репрессиями» (Prigozhin founded the "Foundation for Combating Repression").

50 WHOIS Record for FondfBr.ru

51 DomainTools Investigations (DTI). (2025, March 18). Domain Registrars Powering Russian Disinformation: A Deep Dive into Tactics and Trends. (Retrieved June 30, 2025)

52 Less info about private registrants visible in Whois

53 Wickramasinghe, Nimesha & Nabeel, Mohamed & Thilakaratne, Kenneth & Keppitiyagama, Chamath & De Zoysa, Kasun. (2021). Uncovering IP Address Hosting Types Behind Malicious Websites. (Retrieved July 11, 2025)

54 Why 90 days

55 Council of the European Union. (2022, March 2). EU imposes sanctions on state-owned outlets RT/Russia Today and Sputnik's broadcasting in the EU. Council of the European Union Press Release. Retrieved May 7, 2025.

56 Council of the European Union. (2022, December 16). Council Decision (CFSP) 2022/2478 of 16 December 2022 amending Decision 2014/512/CFSP concerning restrictive measures in view of Russia's actions destabilising the situation in Ukraine. *Official Journal of the European Union*, L 322I, 614–686. Retrieved May 7, 2025.

57 Balint, K., Wildon, J., Arcostanzo, F., & Reyes, K. D. (2022, October 6). Effectiveness of the sanctions on Russian state-affiliated media in the EU: An investigation into website traffic & possible circumvention methods. Institute for Strategic Dialogue. Retrieved May 7, 2025.

58 an open-source intelligence and forensics program

59 The TGStat page for "The World Today with Yuriy Podolyaka"

60 ERR is a metric that measures user interaction (likes, comments, shares) as a proportion of reach. For a channel of this size, such elevated ERR figures are rare: ERRs typically decrease as subscriber numbers grow.

61 PR-CY "Analytics of Telegram channels"

62 Forwarding data and mention graphs provided by TGStat for "The World Today with Yuriy Podolyaka"

63 The Roskomnadzor list of registered personal pages and channels

64 The analysis was done using the Osavul platform, on of Ukraine's situational awareness and attribution systems.

65 Center for Strategic Communications and Information Security. (2024, April 16). "How Russian propaganda speculates on the topic of corruption in Ukraine". (Retrieved June 16, 2025)

66 "Digital Army of Russia" Telegram channel

67 An example of a "task" with multilingual comment templates and dissemination instructions posted in the "Digital Army of Russia" Telegram channel.

68 An example of a "task" with multilingual comment templates and dissemination instructions posted in the "Digital Army of Russia" Telegram channel.

69 Pro-Kremlin Telegram channels promote narrative that Poland will annex western Ukraine | by @DFRLab; Osadchuk, Gigitashvili, 06 May 2022.

70 Russian fake in Warsaw Metro: "Stand up for protection of ancestral Polish lands", 2 February 2023.

71 DFRLab (6 May 2022). Pro-Kremlin Telegram channels promote narrative that Poland will annex western Ukraine.

72 Russian War Report: New fires and alleged sabotage operations across Russian territory; Osadchuk, 03 May 2022.

73 Osadchuk, R., & Gigitashvili, G. (2022, May 6). Pro-Kremlin Telegram channels promote narrative that Poland will annex western Ukraine. Digital Forensic Research Lab.

74 Ukraine war: False claims spread about military movements in Poland and Finland; Holroyd, 06 May 2022.

75 Center for Strategic Communications and Information Security. (2024, April 16). "How Russian propaganda speculates on the topic of corruption in Ukraine". (Retrieved June 16, 2025)

76 An incident is a group of messages on the same topic that promote a certain opinion or similarly describe a single informational event. Incidents that had a minimum of 6 messages in their group have been used for this research. 86 incidents aimed at promoting the topic of corruption in Ukraine and containing messages with elements of threatening tactics have been recorded.

77 Brave UA. (n.d.). Bravery to be Ukraine. Retrieved April 9, 2025.

78 Puente. (2023, October 7). No! Questa ricevuta Cartier da 1 milione di dollari non può essere di Olena Zelenska. Open. Retrieved April 9, 2025.

79 The analysis was conducted using Osavul, an artificial intelligence platform designed to monitor information environments and detect potential threats.

80 Center for Strategic Communications and Information Security. (2024, April 16). "How Russian propaganda speculates on the topic of corruption in Ukraine". (Retrieved June 16, 2025)

81 Kyiv International Institute of Sociology. (2023, November 1). "Public perception of the main problems (except war) and who should make efforts to fight corruption: results of a telephone survey conducted September 30-October 11, 2023". (Retrieved July 14, 2025)

82 The Guardian. (2023, June 20). West needs strategy to tie Ukraine aid to corruption progress, thinktank says. (Retrieved July 14, 2025)

83 Centre for Strategic Communications and Information Security. (2024, November 20). "Exploiting TikTok for malicious influence on ukrainian audience". (Retrieved July 14, 2025)

84 The Territorial Centre of Recruitment and Social Support (TRC or ТЦК) is Ukraine's military administration body that keeps military records and mobilises the population

85 OPORA. (2024, July 10). Media Consumption of Ukrainians: the Third Year of a Full-scale War. (Retrieved July 15, 2025)

86 OPORA. (2024, July 10). Media Consumption of Ukrainians: the Third Year of a Full-scale War. (Retrieved July 15, 2025)

Prepared and published by the
NATO STRATEGIC COMMUNICATIONS
CENTRE OF EXCELLENCE

The NATO Strategic Communications Centre of Excellence (NATO StratCom COE) is a NATO accredited multi-national organisation that conducts research, publishes studies, and provides strategic communications training for government and military personnel. Our mission is to make a positive contribution to Alliance's understanding of strategic communications and to facilitate accurate, appropriate, and timely communication among its members as objectives and roles emerge and evolve in the rapidly changing information environment.

www.stratcomcoe.org | @stratcomcoe | info@stratcomcoe.org